Robust and adaptive search

Yann Disser* Stefan Kratsch†

February 21, 2017

arXiv:1702.05932v1 [cs.DS] 20 Feb 2017

Binary search finds a given element in a sorted array with an optimal number of $\log n$ queries. However, binary search fails even when the array is only slightly disordered or access to its elements is subject to errors. We study the worst-case query complexity of search algorithms that are robust to imprecise queries and that adapt to perturbations of the order of the elements. We give (almost) tight results for various parameters that quantify query errors and that measure array disorder. In particular, we exhibit settings where query complexities of $\log n+ck$, $(1+\varepsilon)\log n+ck$, and $\sqrt{cnk}+o(nk)$ are best-possible for parameter value k, any $\varepsilon > 0$, and constant c.

1. Introduction

Imagine a large register with n files from which you wish to extract a particular file. All files are indexed by some key and the files are sorted by key value. Not knowing the distribution of the keys, you probably use binary search since looking at $\log n$ keys is best possible in the worst case. Unfortunately, however, other users have accessed files before you and have only returned the files to approximately the right place. As a result, the register is unsorted, but at least each file is within some small number k of positions of where it should be. How should you proceed? If you knew k and n, at what ratio of k vs. n should you resort to a linear search of the register? If you do not know k, can you still do reasonably well? What if the register was recently moved, by packing the files into boxes, but in the process the order of the boxes got mixed up, and now there are large blocks of files that are far away from their correct locations? What if you misread some of the keys? Situations like these are close to searching in a sorted register and there are plenty of parameters that measure closeness to a sorted array, e.g., maximum displacement or minimum block moves to sort, respectively persistent or temporary read errors. We give (almost) optimal algorithms for a large variety of these measures, and thereby establish for each of them exact regimes in which we can outperform a linear search of all elements, or even be almost as good as binary search.

More formally, we study the fundamental topic of comparison-based search, which is central to many algorithms and data structures [20, 24, 31]. In its most basic form, the search problem can be phrased in terms of locating an element e within a given array A. In order to search A efficiently, we need structure in the ordering of its elements: In general, we cannot hope to avoid querying all entries to find e. The most prominent example of an efficient search algorithm that exploits special structure is *binary search* for sorted arrays. Binary search is best-possible for

*Main work done while at TU Berlin. Current address: TU Darmstadt, disser@mathematik.tu-darmstadt.de
†University of Bonn, kratsch@cs.uni-bonn.de

this case. It needs only logarithmically many queries and is thus very well suited for searching extremely large collections of data. However, it heavily relies on perfect order and reliable access to the data. For large and dynamically changing collections of data, both requirements may be difficult to ensure, but it may be reasonable to assume the number of imperfections to be bounded. Accordingly, we ask: *What is the best-possible search algorithm if the data may be disordered or we cannot access it reliably? In what regime of the considered measure is it better than linear search?*

We provide (almost) tight bounds on the query complexity of searching an array A with n entries for an element e in a variety of settings. Each setting is characterized by bounding a different parameter k that quantifies the imperfections regarding either our access to array elements or regarding the overall disorder of the data. Note that one can always resort to linear search, which rules out lower bounds stronger than n comparisons.[†] Table 1 gives an overview of the parameters we analyze and our respective results. Qualitatively, our results can be grouped into three groups of settings leading to different query complexities, and we briefly highlight each group in the following.

The first group contains the parameters k_{sum}, $k_{\max}$, and k_{inv}, which quantify the summed/maximum distance of each element from its position in sorted order and the number of element pairs in the wrong relative order, respectively (detailed definitions can be found the the corresponding sections). For all of these parameters we are able to show that $\log n + ck$ queries are necessary and sufficient, for constant c. Intuitively, this is the best complexity we can hope for: We cannot do better than $\log(n)$ queries, and the impact of k on the query complexity is linear and can be isolated.

The second group of results is with respect to the parameters k_{lies}, k_{faults}, as well as multiple parameters for edit distances that measure the number of element operations needed to sort A. The parameter k_{lies} limits the number of queries that yield the wrong result, and k_{faults} limits the number of array positions that yield wrong query outcomes. For bounded values of k_{lies} and k_{faults} we show that e cannot be found with $\log n + ck$ queries using any binary-search-like algorithm.[1] On the other hand, we provide an algorithm that needs $(1 + 1/c) \log n + ck$ queries, for any $c \geq 1$. For bounded edit distances, it is easy to see that we need n queries if e need not be at its correct position relative to sorted order, since e can be moved anywhere with just 2 edits, forcing us to scan the whole array. If we assume e to be at its correct location, we can carry over the results for k_{lies} and k_{faults} to obtain the same bounds for the edit-distance related parameters k_{rep}, k_{seq}, k_{mov}, and k_{swap}.

Lastly, we consider the parameter k_{ainv} that counts the number of adjacent elements that are in the wrong relative order, as well as several parameters measuring the number of block operations needed to sort A. Intuitively, these settings are much more difficult for a search algorithm, as it takes relatively small parameter values to introduce considerable disorder. For the case that e is guaranteed to be at the correct position, we show that $\sqrt{cnk} + o(nk)$ queries are necessary and sufficient to locate e.

The algorithms for k_{ainv} and related parameters assume that the parameter value is known to the algorithm a priori. In contrast, all our other algorithms are oblivious to the parameter, in the sense that they do not require knowledge of the parameter value as long as the target element e is guaranteed to be present in the array. Note that if e need not be present and we have no bound on the disorder, we generally need to inspect every entry of the array in case

[†]Accordingly, all (lower) bounds of the form $f(n, k)$ throughout the paper are to be understood as $\min\{f(n, k), n\}$. A naive bound of n can easily be obtained by scanning the whole array.

[1]We interpret the array as a binary tree (rooted at entry $n/2$, with the two children $n/4$, $3n/4$, etc.), and call an algorithm "binary-search-like" if it never queries a node (other than the root) before querying its parent.

we cannot find e. For the parameter k_{lies}, we do not even know how long we need to continue querying the same elements until we may conclude that e is not part of the array. Any of our oblivious algorithms can trade the guarantee that $e \in A$ against knowledge of the parameter value k: Compute from k the maximum number m of queries that it would take without knowing k when $e \in A$. If the algorithm does not stop within m queries then it is safe to answer that e is not in A.

Overall, our results point out several parameters for which a fairly large regime of k (as a function of n) allows search algorithms that are provably better than linear search. For example, while moving only a single element by a lot can lead to bounds of $\Omega(n)$ on the values of several parameters, and hence trivial guarantees, moving many elements by at most k places gives $k_{\max} = k$ and yields better bounds than linear search (roughly) for $k < \frac{n}{3}$, and as good as binary search when $k = \mathcal{O}(\log n)$. Moving only few elements by an arbitrary number of spaces, in turn, still leads to good bounds via parameters such as k_{mov} or k_{swap}, as long as the target is in the correct place. Parameters such as k_{ainv} grow even more slowly, for certain types of disorder, but, on the other hand, only a small regime allows for better than trivial guarantees. While, for each individual parameter we study, there are "easily searchable" instances where the parameter becomes large and makes the corresponding bound trivial, our results often allow for good bounds by resorting to a different parameter.

1.1. Related Work

Our work falls into the area of *adaptive analysis of algorithms*, which aims at a fine-grained analysis of polynomial-time algorithms with respect to structural parameters of the input. An objective of this field is to find algorithms whose running-time dependence on input size and the structural parameters interpolates smoothly between known (good) bounds for special cases and the worst-case bound for general inputs. The topic of adaptive sorting, i.e., sorting arrays that are presorted in some sense, has attracted a lot of attention, see, e.g., [4, 13, 23, 28].

We now discuss results that are specific to searching in arrays. Several authors addressed the question of how much preprocessing, i.e., sorting, helps for searching, if we take into account the total time investment [8, 22, 29]. Fredman [18] gave lower bounds on searching regarding both queries and memory accesses. A classic work of Yao [32] established that the best way of storing n elements in a table such as to minimize number of queries for accessing an element is by keeping the elements sorted, which requires $\log n$ queries, provided that the key space is large enough. Regarding searching in (partially) unordered arrays, there is a nice result of Biedl et al. [5] about insertion sort based on repeated binary searches.

Under appropriate assumptions, namely that array is sorted and its elements are drawn from a known distribution (e.g., searching for a name in a telephone book), one can do much better than binary search, since the distribution allows a good prediction of where the target should be located. In this case $\mathcal{O}(\log \log n)$ queries suffice on average (cf. [31]); to avoid having to query the entire array, previous work suggests combinations of algorithms that perform no worse than binary search in the worst case [10, 6]. Another interesting branch of study is related to search in arrays of more complicated objects such as (long) strings [1, 17] or abstract objects with nonuniform comparison cost [19, 2].

Many papers have studied searching in the presence of different types of errors, e.g., [7, 15, 16, 25], see [11, 27] for surveys. A popular error model for searching allows for a linear number of lies [3, 7, 12, 14, 26], for which Borgstrom and Kosaraju [7] gave an $\mathcal{O}(\log n)$ search algorithm. In constrast, we bound the number of lies separately via the parameter k_{lies}. Rivest et al. [30] gave an upper bound of $\log n + k \log \log n + \mathcal{O}(k \log k)$ queries for this parameter. Their algorithm

Table 1: Overview of our results, with main results in boldface.† (^{o}: even if oblivious to parameter value; ^{c}: for all $c \geq 1$; ^{t}: for tree-algorithms; ^{e}: for $\mathrm{pos}(e) = \mathrm{rank}(e)$)

		bounds	
parameter	description	lower	upper
Section 3 – number of imprecise queries			
k_{lies}	wrong outcomes	$\boldsymbol{\log n + ck}$ [Th. 3]^{ct}	$\boldsymbol{(1+\frac{1}{c})\log n + (2c+2)k}$ [Th. 2]^{oc}
k_{faults}	indices with wrong outcomes	$\log n + ck$ [Co. 1]^{ct}	$(1+\frac{1}{c})\log n + (2c+2)k$ [Th. 4]^{oc}
Section 4.1 – displacement of elements			
k_{sum}	total displacement	$\boldsymbol{\log n/k + 2k + \mathcal{O}(1)}$ [Th. 5,6]^{o}	
k_{max}	maximum displacement	$\boldsymbol{\log n/k + 3k + \mathcal{O}(1)}$ [Th. 7,8]^{o}	
Section 4.2 – number of inversions			
k_{inv}	all inversions	$\log n/k+2k+\mathcal{O}(1)$ [Co. 3]	$\log n/k + 4k + \mathcal{O}(1)$ [Co. 3]^{o}
k_{ainv}	adjacent inversions	$\boldsymbol{\sqrt{8nk} + o(\sqrt{nk})}$ [Th. 9,10]^{e}	
Section 4.3 – element operations needed to sort the array			
k_{rep}	element replacements	$\log n + ck$ [Co. 4]^{cte}	$(1+\frac{1}{c})\log n + (4c+4)k$ [Th. 11]^{oe}
k_{seq}	$n - \lvert\text{max ordered subseq.}\rvert$	$\log n + ck$ [Co. 4]^{cte}	$(1+\frac{1}{c})\log n + (4c+4)k$ [Th. 11]^{oe}
k_{mov}	element moves	$\log n + ck$ [Co. 4]^{cte}	$(1+\frac{1}{c})\log n + (4c+4)k$ [Th. 11]^{oe}
k_{swap}	element swaps	$\log n + ck$ [Co. 4]^{cte}	$(1+\frac{1}{c})\log n + (8c+8)k$ [Th. 11]^{oe}
k_{aswap}	adj. element swaps	$\log n/k+2k+\mathcal{O}(1)$ [Co. 5]	$\log n/k + 4k + \mathcal{O}(1)$ [Co. 5]^{o}
Section 4.4 – block operations needed to sort the array			
k_{bswap}	block swaps	$4\sqrt{nk} + o(\sqrt{nk})$ [Co. 7] [Th. 13]^{e}	
k_{rbswap}	equal size block swaps	$2\sqrt{2nk} + o(\sqrt{nk})$ [Co. 8]^{e}	$4\sqrt{nk} + o(\sqrt{nk})$ [Th. 12]^{e}
k_{bmov}	block moves	$2\sqrt{2nk} + o(\sqrt{nk})$ [Th. 14]^{e} [Co. 6]	

is based on a continuous strategy for the (equivalent) problem of finding an unknown value in $[1, n]$, upto a given precision, using few yes-no questions. Our algorithm (Theorem 2) uses asymptotically fewer queries if $k_{\text{lies}} = \omega(\log n / \log\log n)$.[2]

The works of Finocchi and Italiano [16] and Finocchi et al. [15] consider a parameter very similar to k_{faults}, with the additional assumption that faults may affect also the working memory of the algorithm, except for $\mathcal{O}(1)$ "safe" memory words. Finocchi and Italiano [16] give a deterministic searching algorithm that needs $\mathcal{O}(\log n + k^2)$ queries. Brodal et al. [9] improve this bound to $\mathcal{O}(\log n + k)$ and Finocchi et al. [15] provide a lower bound of $\Omega(\log n + k)$ even for randomized algorithms. Our results are incomparable as our result for parameter k_{faults} uses only $(1+\frac{1}{c})\log n+(2c+2)k$ queries, getting arbitrarily close to $\log n+O(k)$ (cf. Theorem 4), but does not consider faults in the working memory; the high level approach of balancing progress in the search with security queries is the same as in [9], but more careful counting is needed to get small constants. For parameter k_{lies} we give a simpler algorithm with $2\log n + 4k$ queries and using only $\mathcal{O}(1)$ words of working memory, but it is not clear whether the result can be transferred to k_{faults} without increasing the memory usage.

[2] A technical report of Long [21] claims that the actual tight bound of the algorithm of Rivest et al. [30] is $\mathcal{O}(\log n + k)$, which is consistent with our results.

Finally, we comment on the measures of disorder we adopt in this paper. We study various well-known measures that are mostly folklore. Detailed overviews of measures and their relations were given by Petersson and Moffat [28] and Estivill-Castro and Wood [13]. For the sake of completeness and to get all involved coefficients, we have provided in Appendix A proofs of all pairwise relations between our parameters, as depicted in Figure 1.

2. Preliminaries

In this paper we consider the following problem: Given an array A of length n and an element e, find the position of e in A or report that $e \notin A$ with as few *queries* as possible. We use $A[i]$, $i \in 1, \ldots, n$ to denote the i-th entry of A. We allow access to the entries of A only via queries to its indices, regarding the relation of the corresponding element to e. We write query(i) for the operation of querying A at index i, and let query$(i) =$'¡' (respectively, '¿' or '=') denote the outcome indicating that $A[i] < e$ (respectively $A[i] > e$ or $A[i] = e$). Note that in faulty settings the query outcome need not be accurate.

To keep notation simple, we generally assume the entries of A to be unique unless explicitly stated otherwise. We emphasize that none of our results relies on this assumption. We can then define pos(a) to denote the index of a in A, by setting pos$(a) = i$ if and only if $A[i] = a$. Further, let rank$(a) = |\{i : A[i] < a\}| + 1$ be the "correct" position of a with respect to a sorted copy of A, irrespective of whether or not $a \in A$. We often use an element $a \in A$ and its index pos(a) interchangeably, especially for the target element e. Note that, as discussed in the introduction, for oblivious algorithms we generally assume $e \in A$.

3. Searching with imprecise queries

In this section, we consider the problem of finding the index pos(e) of an element e in a sorted array A of length $n = 2^d$, $d \in \mathbb{N}$ in a setting where queries may yield erroneous results. We say that '¡' is a lie (the truth) for index i if $A[i] \geq e$ ($A[i] < e$), and analogously for '¿' and '='. To quantify the number of lies, we introduce two parameters k_{lies} and k_{faults}. The first parameter k_{lies} simply bounds the number of queries with erroneous results, which we interpret as the number of lies allowed to an adversary. The second parameter k_{faults} bounds the number of indices i for which query(i) (consistently) returns the wrong result, allowing the conclusion that $e \notin A$ in case query(e) yields the wrong result. Equivalently, for an unsorted array A, we can require all queries to be truthful and define $k_{\text{faults}}(e)$ to be the number of inversions involving e, i.e., $k_{\text{faults}}(e) = |i : (i < \text{pos}(e) \wedge A[i] > e) \vee (i > \text{pos}(e) \wedge A[i] < e)|$. Observe that both definitions of k_{faults} are equivalent. For clarity, we write k_{faults} when considering the adversarial interpretation, and $k_{\text{faults}}(e)$ when considering it as a measure of disorder of an unsorted array. For both k_{lies} and k_{faults}, we only allow queries to e to yield '='.

The algorithms of this section operate on the binary search tree rooted at index $r = n/2$ that contains a path for each possible sequence of queries in a binary search of the array, and identify nodes of the tree with their corresponding indices. We write $\text{next}_{>}(i)$ and $\text{next}_{<}(i)$ to denote the two successors of node i, e.g., $\text{next}_{>}(r) = n/4$ and $\text{next}_{<}(r) = 3n/4$. Similarly, we write prev(i) to denote the predecessor of i in the binary search tree, and $\text{prev}_q(i) = v$ for the last vertex v on the unique r-i-path such that $\text{next}_q(v)$ also lies on the r-i-path ($\text{prev}_q(i) = \emptyset$ if no such node exists). Intuitively, $\text{prev}_q(i)$ is the last vertex corresponding to an array entry larger (if $q =$"¿") or smaller (if $q =$"¡") than $A[i]$. For convenience, query$(\emptyset) = \emptyset$, $\text{prev}_{<>}(r) = \emptyset$, and $\text{next}_{<>}(i) = i$ if i is a leaf of the tree. We further denote by $d(i, j)$ the length of the path from

node i to node j in the search tree. We say that an algorithm *operates on the binary search tree* if no index is queried before its predecessor in the tree.

We start by considering the parameter k_{lies}. If we knew the value of this parameter, we could try a regular binary search, replace every query with $2k_{\text{lies}} + 1$ queries to the same element and use the majority outcome in each step. However, this would give $(2k_{\text{lies}} + 1) \log n$ queries, where ideally we should not use more than $\log n + f(k)$ queries. We first give an algorithm that achieves the separation between n and k_{lies} while being oblivious to the value of k_{lies}. Importantly, the algorithm only needs $\mathcal{O}(1)$ memory words, which also makes it applicable to settings where "safe" memory, that cannot be corrupted during the course of the algorithm, is limited. This algorithm still needs $2 \log n + f(k)$ queries, but we will show later how to build on the same ideas to (almost) eliminate the factor of 2.

Algorithm 1: Algorithm with $2 \log n + 4k_{\text{lies}}$ queries

```
the algorithm stops once a query yields '='
i ← n/2                                          // start at the root
while (q ← query(i)) ≠ '=' do                    // by definition, '=' cannot be a lie
    i' ← prev_¬q(i)                              // ∅ if all queries on the path from the root yielded q
    while i ≠ i' ∧ query(i') = q do              // while query(i') contradicts its previous outcome...
        i ← prev(i)                              // ...backtrack towards i'
    if i ≠ i' then                               // if we did not backtrack all the way to i'...
        i ← next_q(i)                            // ...proceed according to q
```

Intuitively, Algorithm 1 searches the binary search tree defined above, simply proceeding according to the query outcome at each node. In addition, the algorithm invests queries to double check past decisions. We distinguish left and right turns, depending on whether the algorithm proceeds with the left or the right child. In particular, before proceeding, the algorithm queries the last vertex on the path from the root where it decided for a turn in the opposite direction. While an inconsistency to previous queries is detected, i.e., a query to a vertex where it turned right (or left) gives '¿' (or '¡'), the algorithm backtracks one step. In this manner, the algorithm guarantees that it never proceeds along a wrong path without the adversary investing additional lies. Note that if the algorithm only ever turned right (or left), i.e., there was no previous turn in the opposing direction, it does not double check any past decisions until the query outcome changes. This is alright since either the algorithm is on the right path or the adversary needs to invest a lie in each step.

Theorem 1. *We can find e obliviously using $2 \log n + 4k_{\text{lies}}$ queries and $\mathcal{O}(1)$ memory.*

Proof. We claim that Algorithm 1 achieves the bound of the theorem. Note that $\text{prev}_{\neg q}(i)$ only depends on i and not on the outcome of previous queries, therefore, we can determine it with $\mathcal{O}(1)$ memory words. We will show that in each iteration of the outer loop of the algorithm, the potential function $\Phi = 2d(i, e) + 4k$ decreases by at least one for each query, where k is the number of remaining lies the adversary may make. This proves the claim, since $\Phi \geq 0$ and initially $\Phi \leq 2 \log n + 4k_{\text{lies}}$. We analyze a single iteration of the outer loop.

Observe that if z is the number of iterations of the inner loop, then the total number of queries is $z + 2$ if the inner loop terminates because $\text{query}(i') = \neg q$, and $z + 1$ if it terminates because $i = i'$. If an iteration of the inner loop is caused by $\text{query}(i')$ being a lie, then in this iteration $\Delta\Phi \leq 2 - 4 = -2$, and otherwise, $d(i, e)$ is decreased by one and likewise $\Delta\Phi = -2 + 0 = -2$. Overall, the change in potential during the inner loop is always $\Delta\Phi = -2z$. If the inner loop

terminates because $i = i'$, then $z \geq 1$ and the total change in potential is $\Delta\Phi \leq -2z \leq -z - 1$, enough to cover all $z + 1$ queries.

Now consider the case that the inner loop terminates because $\text{query}(i') = \neg q$. If $\neg q$ is a lie for i' or q is a lie for i, the adversary invested an additional lie, and even if the last update to i increases $d(i, e)$, the total change in potential is bounded by $\Delta\Phi \leq -2z - 4 + 2 \leq -2z - 2$, enough to cover all $z + 2$ queries. On the other hand, if $\neg q$ is the truth for i' and q is the truth for i, then $e \in \{i', \ldots, i\}$ and i must lie on the unique r-e-path in the search tree (and $i \neq e$). The final update to i thus decreases $d(i, e)$ by 1 and the total change in potential is $\Delta\Phi = -2z - 2$, again enough to cover all $z + 2$ queries. □

Algorithm 2: Algorithm with $(1 + \frac{1}{c}) \log n + (2c + 2)k_{\text{lies}}$ queries

```
the algorithm stops once a query yields '='
i ← n/2                                          // start at the root
while (q ← query(i)) ≠ '=' do                    // by definition, '=' cannot be a lie
    i' ← prev_¬q(i)                              // ∅ if all queries on the path from the root yielded q
    while 0 < cΔ_i' < d(i, i') + 1 do            // while we do not have sufficient support to proceed...
        query(i')                                // ...query i' for support
    if Δ_i' = 0 then                             // if we ran out of support at i' altogether...
        i ← i'                                   // ...backtrack to i'
    else                                         // if we have sufficient support at i'...
        i ← next_q(i)                            // ...proceed according to q
```

We now adapt Algorithm 1 to minimize the impact of potential lies on the dependency on $\log n$ in the running time. Intuitively, instead of backing up each query $q \leftarrow \text{query}(i)$ by a query to $\text{prev}_{\neg q}(i)$, we back only one in c queries (cf. Algorithm 2). During the course of the algorithm and its analysis, we let $n_{q,j}$ denote the number of queries (so far) to node j that resulted in $q \in \{<, >\}$ and $\Delta_j := |n_{<,j} - n_{>,j}|$.

Theorem 2. *For every $c \geq 1$, we can find e obliviously using $(1 + \frac{1}{c}) \log n + (2c + 2)k_{\text{lies}}$ queries.*

Proof. We claim that Algorithm 2 achieves the bound of the theorem. In this algorithm, we intuitively back up every c-th query (for integral c). To capture this in our potential function, we need a term that stores potential for the next c queries. We will introduce two such terms L, T, representing the case the algorithm's current belief of the relation between i' and e is a lie or the truth, respectively. We need to distinguish these cases, since they lead to different behavior regarding $d(i, e)$ and the number of remaining lies.

We need the following additional notation. For some current value of i during the execution of the algorithm, we define the type of a node j of the search tree on the r-i-path to be $t_j \in \{<, >\}$ if $\text{next}_{t_j}(j)$ also lies on this path. Further, we let $\text{succ}_q(j) = j'$ if j' is the first node on the j-i-path with $t_{j'} = q$ or $\text{succ}_q(j) = i$ if no such node exists. We set $\text{succ}_q(i) = \emptyset$. To avoid special treatment of leaves, we replace each leaf of the search tree by an infinite binary tree of nodes corresponding to the original leaf, in both algorithm and analysis. If e was a leaf, then, for each new node j corresponding to e, we set $d(j, e) = d(j, r_e)$ where r_e is the root of the subtree corresponding to e.

Intuitively, the potential of the algorithm needs to depend on $c\Delta_{i'} - d(i, i')$, since this difference captures the number of steps it can still make before it needs to use a backup query. To keep track of this difference across iterations of the algorithm, we introduce the notion of a *zig-zag*

pair, which we will define formally below. In particular, (i, i') always forms a zig-zag pair. Let $j = \text{next}_q(i)$ in some iteration after which $\Delta_{i'} \neq 0$, i.e., i gets updated to j. If j has the same type as i in the next iteration, i' stays the same and we can simply replace the zig-zag pair (i, i') with (j, i'). On the other hand, if j has a different type in the next iteration, we need to introduce a new pair (j, i). Since we may backtrack later and continue differently at i, we also need to keep the pair (i, i'). Conceptually, we need to keep track of all maximal l-l'-subpaths of the current r-i-path with the property that l' has the opposite type than all other nodes on the subpath. If the algorithm backtracks to node l at some point, then, in the next iteration, $i = l$ and $i' = l'$, and the difference $c\Delta_{l'} - d(l, l')$ captures how much potential remains to continue querying without using a back up query to $i' = l'$.

Formally, we define the set of all zig-zag pairs as

$$Z := \left\{(j, j') \mid \exists q \in \{<, >\}.j' = \text{prev}_q(j) \wedge j = \text{succ}_q(j')\right\}.$$

Note that $(i, i') \in Z$ throughout the algorithm, and that every node appears at most once as the second element of a zig-zag pair, exactly if it has a different type than its successor on the unique r-i-path. For convenience, we set $\Delta_\emptyset$ to be equal to the number of all "queries" to $\emptyset$ in the algorithm. We define

$$L = \sum_{(j,j') \in Z} [c\Delta_{j'} - d(j, j')] \cdot \Lambda_{t_{j'}, j'},$$

where $\Lambda_{q,j} = 1$ if q is a lie for j and $\Lambda_{q,j} = 0$ otherwise. Similarly, we define

$$T = \sum_{(j,j') \in Z} [d(j, j') - c(\Delta_{j'} - 1)] \cdot (1 - \Lambda_{t_{j'}, j'}).$$

With this notation in place, we introduce the extended potential function

$$\Phi = (1 + \frac{1}{c})d(i, e) + (2 + \frac{1}{c})L + \frac{1}{c}T + (2c + 2)k,$$

where k is the number of lies remaining to the adversary.

We claim that $L, T \geq 0$ and thus $\Phi \geq 0$ holds after each iteration. More precisely, we show that the contribution of each zig-zag pair in Z to either L or T is non-negative. To see this, first observe that in each iteration Z changes exactly by either removing the zig-zag pair (i, i') (unless $i = r$), by replacing it with the pair $(\text{next}_q(i), i')$, or by adding a new pair $(\text{next}_q(i), i)$. Inductively, it thus suffices to show that the contribution of $(\text{next}_q(i), i')$ or $(\text{next}_q(i), i)$ in the latter cases ($\Delta_{i'} \neq 0$) is positive. First, observe that $\Delta_i = 1$ after the iteration, hence, if $(\text{next}_q(i), i) \in Z$, its contribution to L or T must be positive.

Now consider the case that $(\text{next}_q(i), i') \in Z$ after the iteration. By definition of the algorithm, the inner loop ensures that

$$c\Delta_{i'} \geq d(i, i') + 1 = d(\text{next}_q(i), i'),$$

hence the contribution of $(\text{next}_q(i), i')$ to L is non-negative. Now consider the last iteration of the outer loop in which $\Delta_{i'}$ changed, and let j, j' be the corresponding values of i and i' in that iteration. Either $d(i, i') = \Delta_{i'} = 1$, or the last change to $\Delta_{i'}$ was because $j' = i'$ and $c\Delta_{i'} < d(j, i') \leq d(i, i') + 1$. In the latter case, after the update to $\Delta_{i'}$, we have $c\Delta_{i'} \leq d(j, i') + 1 + c$ and thus, in both cases,

$$c\Delta_{i'} \leq d(i, i') + c < d(\text{next}_q(i), i') + c,$$

hence the contribution of $(\text{next}_q(i), i')$ to T is non-negative.

Initially, $T = L = 0$ since $Z = \emptyset$, and $\Phi = (1+\frac{1}{c})\log n + (2c+2)k_{\text{lies}}$. Since $\Phi \geq 0$ throughout, it thus suffices to show that in each iteration of the outer loop Φ decreases by at least one for each query. We consider a single iteration of the outer loop for fixed i, i'.

First consider the case where the inner loop is not executed. In this case, the algorithm makes only a single query, and we need to show that the potential decreases by at least 1. Observe that L cannot increase during the update to i and T may increase by at most 1. If q is a lie for i, then this update increases $d(i,e)$ by 1 and the adversary invested an additional lie, for a change in potential of at most $\Delta\Phi \leq (1+\frac{1}{c}) + 0 + \frac{1}{c} - (2c+2) = -2c - 1 + \frac{2}{c} \leq -1$ (since $c \geq 1$). If $\neg q$ is a lie for i', then $d(i,e)$ increases by 1 and L decreases by 1, for a change in potential of $\Delta\Phi \leq (1+\frac{1}{c}) - (2+\frac{1}{c}) + 0 + 0 = -1$. If q is the truth for i and $\neg q$ is the truth for i', then $e \in \{i, \ldots, i'\}$ and i must lie on the unique r-e-path. The update to i then decreases $d(i,e)$ and changes the potential by $\Delta\Phi = -(1+\frac{1}{c}) + 0 + \frac{1}{c} - 0 = -1$.

Now consider the case where the inner loop is executed until $\Delta_{i'} = 0$, and fix the value of $\Delta_{i'}$ before the inner loop. We may assume that no query to i' yielded $\neg q$, otherwise we can balance each such query with a query that yielded q, one the two being a lie, for a change in potential of $\Delta\Phi = -(2c+2) \leq -4$, which pays for both these queries. With this assumption, we have exactly $\Delta_{i'} \geq 1$ queries in the loop, all of which yielded q. Note that the previous iteration of the outer loop ensured that $c\Delta_{i'} \geq d(i,i')$. If $\neg q$ is a lie for i', the eventual update to i decreases $d(i,e)$ by $d(i,i')$ and decreases L by $c\Delta_{i'} - d(i,i')$ (since (i,i') is eliminated from Z). The overall change in potential then is

$$\begin{aligned}
\Delta\Phi &\leq -(1+\frac{1}{c}) \cdot d(i,i') - (2+\frac{1}{c})[c\Delta_{i'} - d(i,i')] + 0 + 0 \\
&= d(i,i') - c\Delta_{i'} - (c+1)\Delta_{i'} \\
&\leq -(c+1)\Delta_{i'} \\
&\overset{c\geq 1}{\leq} -1 - \Delta_{i'},
\end{aligned}$$

which is enough to cover all $1 + \Delta_{i'}$ queries. On the other hand, if $\neg q$ is the truth for i', the eventual update may increase $d(i,e)$ by at most $d(i,i')$ and it eliminates the contribution of (i,i') to T (since $i = i'$ and, hence, $(i,i') \notin Z$). The adversary invested $\Delta_{i'}$ additional lies, and the change in potential is

$$\begin{aligned}
\Delta\Phi &\leq +(1+\frac{1}{c})d(i,i') + 0 + 0 - (2c+2)\Delta_{i'} \\
&\overset{c\geq 1}{\leq} 2d(i,i') - 4\Delta_{i'} \\
&\leq -2\Delta_{i'} \\
&\leq -1 - \Delta_{i'},
\end{aligned}$$

which is again enough to cover all $1 + \Delta_{i'}$ queries.

Finally, consider the case where the inner loop is executed until $c\Delta_{i'} \geq d(i,i') + 1$. As before, $c\Delta_{i'} \geq d(i,i') \geq 1$, and, again, we may assume that no query to i' yielded q. Hence, as $c \geq 1$, we made a single query to i' that yielded $\neg q$, i.e., 2 queries overall. We need to show that $\Delta\Phi \leq -2$. Assume first that $\neg q$ is the truth for i' and we thus decreased T by $(c-1)$. If q is the truth for i, we have $\Delta\Phi \leq -(1+\frac{1}{c}) + 0 - \frac{1}{c}(c-1) - 0 = -2$. If q is a lie for i, we have $\Delta\Phi \leq (1+\frac{1}{c}) + 0 - \frac{1}{c}(c-1) - (2c+2) = -2 - 2c + \frac{2}{c} \leq -2$. Now assume that $\neg q$ is a lie for i' and we thus increased L by $(c-1)$. Since the adversary invested an additional lie, we have $\Delta\Phi \leq (1+\frac{1}{c}) + (2+\frac{1}{c})(c-1) + 0 - (2c+2) = -2$. □

To provide a strong lower bound, we restrict ourselves to algorithms that operate on the binary search tree. Such algorithms interpret the array as a binary tree (rooted at entry n/2, with the two children n/4, 3n/4, etc.), and never query a node (other than the root) before querying its parent.

Theorem 3. *For every $c \in \mathbb{N}$, no algorithm operating on the search tree can find e with less than $\log n + ck_{\text{lies}}$ queries in general.*[†]

Proof. We consider the behavior of the algorithm on the search tree for large values of n that are powers of two. We split the queries of the algorithm into phases, where phase p starts as soon as a node of depth $(c+1)\cdot p$ is queried for the first time, starting with phase 0. We take the perspective of an adversary and specify the outcome to each query, ensuring that at most one lie is invested in each phase, and at most k_{lies} lies overall. Note that we do not have to decide immediately whether a query outcome is truthful and neither do we have to fix the position of e a priori.

Consider a fixed phase p. The first query of the phase to node i of depth $(c+1)\cdot p$ yields '<', all subsequent queries to positions smaller (larger) than i yield '>' ('<'). If the algorithm queries more than once a node of depth $(c+1)\cdot p$ in the first $c+1$ queries of the phase or any node in the left subtree of i, then the phase needs at least $c+2$ queries and we do not lie, i.e., e is in the subtree rooted at the leftmost node of depth $(c+1)(p+1)$ in the right subtree of i. Otherwise, we lied for the query to node i and e is in the subtree rooted at the rightmost node i' of depth $(c+1)(p+1)$ in the left subtree of i. Since no node in the left subtree of i has been queried yet, the algorithm needs an additional c queries to reach i', for a total of $2c+1$ queries in the phase. Once all lies have been used up, we continue answering queries as before, and each phase trivially needs at least $c+1$ queries.

Observe that querying every node on the path from node $n/2$ to e requires exactly $\log n$ queries, or $c+1$ queries per phase (except maybe a last, partial phase). Now if we use up all k_{lies} lies, then there are k_{lies} phases that need c additional queries each, for a total of $\log n + ck_{\text{lies}}$, as claimed. Otherwise, let $P(n) > \left\lfloor \frac{\log n}{c+1} \right\rfloor - k_{\text{lies}}$ be the number of phases in which we did not lie. Each such phase needed $c+2$ queries instead of $c+1$. Overall, we have more than $\log n + P(n)$ queries. Since $P(n)$ is unbounded with growing n while k_{lies} and c are constant, we have $\log n + P(n) \geq \log n + ck_{\text{lies}}$ for n large enough, as claimed. □

Note that the construction in the proof of Theorem 3 can be applied without change to k_{faults}, since the adversary never gives conflicting replies. As a consequence, we immediately obtain a lower bound for k_{faults}.

Corollary 1. *For every $c \in \mathbb{N}$, no algorithm operating on the search tree can find e with less than $\log n + ck_{\text{faults}}$ queries in general.*

We show how to translate any algorithm with a performance guarantee with respect to k_{lies} to an algorithm with the same guarantee for k_{faults}.

Theorem 4. *Let $f : \mathbb{N}^2 \to \mathbb{N}$. If we can find e with $f(n, k_{\text{lies}})$ queries, then we can find e with $f(n, k_{\text{faults}})$ queries.*

Proof. Assume we have an algorithm that needs $f(n, k_{\text{lies}})$ queries. The difficulty when applying this algorithm for k_{faults} is that, in the faulty setting, there is no benefit in querying the same elements again. However, we can simulate repeated queries to the same element as follows. Say the algorithm needs to query a previously queried element i with the understanding that the

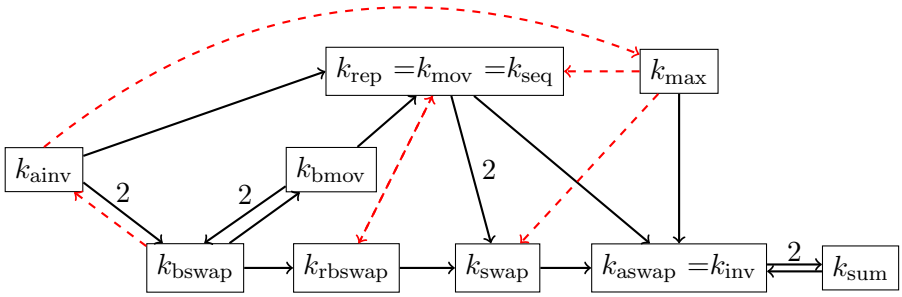

Figure 1: Overview of relations between measures of disorder. A solid black path from k to k' means that $k \leq ck'$, where c is the product of the edge labels along the path ($c = 1$ for unlabeled paths). If there is no solid black path from k to k', then k cannot be bounded by ck' for any constant c. Every arc is proved explicitly in Appendix A (dashed red arcs correspond to unboundedness results), and *all* other relationships are implied.

adversary has to pay for lying repeatedly. Let i' be the first unqueried index to the left or to the right of i. If no such index exists, we already queried all elements and found e, since query(e) is guaranteed to return the correct result. We query i' instead of i. If $i' = e$, we are done. Otherwise, we know that no index in $[i, i']$ contains e, and, hence, all these elements are left of e or all of them are right of e. Therefore the query to i' is equivalent to another query to i when the adversary has to pay for repeated lies. Every fault can be treated as a lie by the adversary, and we get the claimed bound. □

4. Searching disordered arrays

In this section, we consider the problem of finding the index $\text{pos}(e)$ of an element e in array A of length $n = 2^d$, $d \in \mathbb{N}$. In contrast to Section 3, we do not assume A to be sorted but expect all queries to yield correct results. We study a variety of parameters that quantify the disorder of A and provide algorithms and lower bounds with respect to the different parameters. Figure 1 gives the relationship between every pair of parameters. The proofs of these relationships can be found in Appendix A.

4.1. Bounded displacement

We now consider the two parameters k_{sum} and $k_{\max}$ that quantify the displacement of elements between A and $A^\star$. More precisely, we define $k_{\text{sum}} := \sum_{x \in A} |\text{pos}(x) - \text{rank}(x)|$ and $k_{\max} := \max_{x \in A} |\text{pos}(x) - \text{rank}(x)|$. We first derive bounds in terms of k_{sum}.

Theorem 5. *Every search algorithm needs at least* $\lfloor \log(n/2k_{\text{sum}}) \rfloor + 2k_{\text{sum}} + 1$ *queries, even if the elements other than* e *are in the correct relative order.*[†]

Proof. We give a strategy for an adversary to position the elements of the array adaptively, depending on the queries of the search algorithm. The strategy maintains a range $\{l, \ldots, r\}$ of candidate indices for the searched element e that never grows during the course of the strategy. In the beginning, we set $l = 1$ and $r = n$.

In the first phase of the strategy, we maintain the invariant that all queries to indices $i < l$ yield (and yielded) the result $A[i] < e$, and all queries to indices $j > r$ yield $A[j] > e$. Whenever an

index $i \in \{l, \dots, r\}$ is queried, the result depends on whether $\{l, \dots, i\}$ is larger than $\{i, \dots, r\}$ or not. In the former case, the query yields $A[i] > e$ and we set $r = i - 1$. In the latter case, it yields $A[i] < e$ and we set $l = i + 1$. The first phase ends after $\lfloor \log(n/2k_{\text{sum}}) \rfloor - 1 = \lfloor \log(n/4k_{\text{sum}}) \rfloor \leq \lfloor \log(n/(2k_{\text{sum}} + 2)) \rfloor$ queries. At this point, we have $r - l + 1 \geq 2k_{\text{sum}} + 2$, hence there are still at least $2k_{\text{sum}} + 2$ positions left that may contain e.

In the second phase of the adversarial strategy, we answer the next $2k_{\text{sum}} + 1$ queries to indices $i \in \{l, \dots, \lfloor (l + r)/2 \rfloor\}$ with $A[i] < e$ and all queries to $i \in \{\lfloor (l + r)/2 \rfloor + 1, \dots, r\}$ with $A[i] > e$. Afterwards, at least one unqueried index in $\{l, \dots, r\}$ remains. It is easy to see that e being in this position is consistent with all queries so far. Overall, the position of e cannot be found with fewer than $\lfloor \log(n/2k_{\text{sum}}) \rfloor + 2k_{\text{sum}} + 1$ queries, as claimed. Moreover, all $<$ answers are left of $>$ answers, allowing all elements other than e to be in correct relative order. □

We extract the following corollary from the proof of Theorem 5.

Corollary 2. *There is a constant $c \in \mathbb{N}$, such that for every $l > 0$, the adversary can ensure that after $\log n/l + c$ queries, an unqueried subarray of length l remains, such that all elements to the left of the subarray are smaller than e, while all elements to the right of it are larger than e.*

We give an algorithm that achieves a optimal number of queries up to an additive up to an additive gap of $\log k_{\text{sum}} + \mathcal{O}(1)$, while being oblivious of the value of k_{sum}.

Theorem 6. *We can find e obliviously using $\log n/k_{\text{sum}} + 2k_{\text{sum}} + \mathcal{O}(1)$ queries.*

Proof. We first perform a regular binary search for e, ignoring the fact that we may be misguided by elements being displaced. In $\log n + \mathcal{O}(1)$ steps, we find e or an index i with $A[i] < e$ and $A[i + 1] > e$. Let $\Delta_i := |\operatorname{pos}(A[i]) - \operatorname{rank}(A[i])|$. We have $\operatorname{rank}(e) > \operatorname{rank}(A[i]) \geq i - \Delta_i$ and $\operatorname{rank}(e) < \operatorname{rank}(A[i + 1]) \leq i + 1 + \Delta_{i+1}$. With $\operatorname{pos}(e) \in \{\operatorname{rank}(e) - \Delta_e, \dots, \operatorname{rank}(e) + \Delta_e\}$ and $\Delta_e + \Delta_i + \Delta_{i+1} \leq k_{\text{sum}}$, we get

$$\operatorname{pos}(e) \in \{i - k_{\text{sum}} + 1, \dots, i + k_{\text{sum}}\}.$$

We can search this range obliviously by querying the elements $i, i+1, i-1, i+2, i-2, i+3, \dots$ in this order, until we find e. During the initial binary search, we already queried $\log k_{\text{sum}} + \mathcal{O}(1)$ of these elements, hence we need a total number of queries equal to

$$\log n + 2k_{\text{sum}} - \log k_{\text{sum}} + \mathcal{O}(1) = \log n/k_{\text{sum}} + 2k_{\text{sum}} + \mathcal{O}(1).$$

□

We now turn our attention to the parameter $k_{\max}$.

Theorem 7. *Every search algorithm needs at least $\log n/k_{\max} + 3k_{\max} + \mathcal{O}(1)$ queries.*[†]

Proof. By Corollary 2, the adversary can ensure without creating inversions that after using $\log n/4k_{\max} + \mathcal{O}(1) = \log n/k_{\max} + \mathcal{O}(1)$ queries the element e may still be at any position of an unqueried subarray of length $4k_{\max}$. It is therefore sufficient to show that finding e in an array of length $4k_{\max}$ and with maximum displacement $k_{\max}$ may take $3k_{\max} - 2$ queries.

We devise a strategy for the adversary that is split in two phases. In the first phase, we maintain that all queried positions in $L = \{1, \dots, 2k_{\max}\}$ contain elements smaller than e, while the positions in $R = \{2k_{\max} + 1, \dots, 4k_{\max}\}$ contain elements larger than e. The first phase ends when $k_{\max} - 1$ positions have been queried in each half of the array.

Without loss of generality, assume that exactly $k_{\max} - 1$ positions have been queried in L at the beginning of the second phase. Otherwise, this is true for R and the argument proceeds analogously. We now restrict the position of e to definitely lie in L. The second phase proceeds until another $k_{\max} - 1$ positions in L have been queried. All queries to positions in $R \cup \{1\}$ are answered as before. For the queries to positions in $L \setminus \{1\}$ we return the inverse answer to before.

The third phase proceeds until one more position in L is queried, which will contain e.

The number of queries up to this point is at least $2(k_{\max}-1)$ for the first phase, at least $k_{\max}-1$ for the second phase, and at least 1 for the third phase. Hence, the total number of queries is $3k_{\max} - 2$ as claimed. It remains to argue that we can rearrange the array such that all elements smaller (larger) than e are on the left (right) of e while not moving elements by more than $k_{\max}$ positions. We fix a final ordering by requiring that the smaller (larger) elements remain in the same relative order.

If position 1 was not queried in phase 2, then there are $k_{\max} - 1$ elements smaller than e and $k_{\max} - 1$ elements larger than e in L. We have $\text{rank}(e) = k_{\max}$, thus e is displaced by at most $k_{\max}$ positions. All elements smaller than e are displaced by at most $k_{\max}$ positions, since there are at most $k_{\max}$ elements that are greater or equal to e in L. Similarly, all elements larger than e are displaced by at most $k_{\max}$ positions, since there are at most $k_{\max} - 1$ elements that are smaller or equal to e in $L \setminus \{1\}$.

If position 1 was queried in phase 2, then there are $k_{\max}$ elements smaller than e and $k_{\max} - 2$ elements larger than e in L. The element in position 1 has rank 1 and all other elements are displaced by at most $k_{\max}$ positions in $L \setminus \{1\}$, as before. □

To obtain a tight upper bound, we need the following observations.

Proposition 1. *If $A[i] > A[j]$ then $i \geq j - (2k_{\max} - 1)$.*

Proof. If $A[i] > A[j]$ then $\text{rank}(A[i]) \geq \text{rank}(A[j]) + 1$ by definition. Using that $i \geq \text{rank}(A[i]) - k_{\max}$ and $\text{rank}(A[j]) \geq j - k_{\max}$, both by assumption, we derive

$$i \geq \text{rank}(A[i]) - k_{\max} \geq \text{rank}(A[j]) + 1 - k_{\max} \geq j - (2k_{\max} - 1),$$

as claimed. □

Lemma 1. *For all i we have $|\{j < i : A[j] > A[i]\}| \leq k_{\max}$ and symmetrically $|\{j > i : A[j] < A[i]\}| \leq k_{\max}$.*

Proof. For the sake of contradiction, assume that there are $j_1 < j_2 < \cdots < j_{k_{\max}+1} = i - 1$ with $A[j_l] > A[i]$ for all $l \in \{1, 2, \ldots, k_{\max} + 1\}$ (the symmetrical case can be proven analogously). Let $A[r]$ be such that $\text{rank}(A[r]) = j_{k_{\max}+1} - k_{\max} = i - k_{\max} - 1$. We have $r < i$, since $i - \text{rank}(A[r]) = k_{\max} + 1 > k_{\max}$. Also $A[r] < A[j]$ for all $j \geq i$, since $\text{rank}(A[j]) \geq j - k_{\max} \geq i - k_{\max} > \text{rank}(A[r])$. On the other hand, the number of elements $A[j]$ with $j < i$ and $A[j] \leq A[i]$ is at most $j_{k_{\max}+1} - (k_{\max} + 1) = i - k_{\max} - 2$, since the number of elements $A[j]$ with $j < i$ and $A[j] > A[i]$ is at least $k_{\max} + 1$ by assumption. But then, the total number of elements that are smaller than $A[r] < A[i]$ is at most $i - k_{\max} - 3$. This is a contradiction with $\text{rank}(A[r]) = i - k_{\max} - 1$. □

We now describe an algorithm that achieves an optimal number of queries up to an additive constant, while being oblivious to the value of $k_{\max}$.

Theorem 8. *We can find e obliviously using $\log n / k_{\max} + 3k_{\max} + \mathcal{O}(1)$ queries.*

Proof. We first use a binary search to find e or an index i with $a_i < e < a_{i+1}$ with $\log n + \mathcal{O}(1)$ queries. By Proposition 1, we have $\mathrm{pos}(e) \in W = \{i - 2k_{\max} + 1, \dots, i + 2k_{\max}\}$. We query the positions in W, starting from the center (positions i and $i+1$ don't need to be queried again) and moving to the left whenever the number of queried elements of W larger than e exceeds the number of smaller elements, and moving to the right otherwise. This can be done obliviously, i.e., without knowing $k_{\max}$.

We claim that we are guaranteed to encounter e within $3k+1$ queries. To see this, assume without loss of generality that e is in the left half of W. In this case, we claim that we do not query any elements in $\{i + k_{\max} + 2, \dots, i + 2k_{\max}\}$. For the sake of contradiction, assume we query positions $\{i+1, \dots, i + k_{\max} + 2\}$, i.e., at least $k_{\max} + 2$ positions in the right half of W. This means that we queried at least $k_{\max} + 1$ elements smaller than e in W before finding e, by construction of the algorithm. But then $|\{j > \mathrm{pos}(e) : A[j] < e\}| > k_{\max}$, contradicting Lemma 1.

We can refine this analysis by observing that we already queried at least $\log k + \mathcal{O}(1)$ positions among $\{i - k_{\max} + 1, \dots, i + k_{\max}\}$ during the initial binary search. □

4.2. Inversions

In this section we consider the number of inversions between elements of the array A. More precisely, we define the number of inversions to be $k_{\mathrm{inv}} := |\{i < j : A[i] > A[j]\}|$, and the number of adjacent inversions to be $k_{\mathrm{ainv}} := |\{i : A[i] > A[i+1]\}|$.

We have $k_{\mathrm{sum}} \le k_{\mathrm{inv}} \le 2k_{\mathrm{sum}}$ (cf. Proposition 13), therefore the results for k_{sum} (Theorems 5 and 6) carry over to k_{inv} with a gap of 2.

Corollary 3. *Every search algorithm needs at least* $\log n/k_{\mathrm{inv}} + 2k_{\mathrm{inv}} + \mathcal{O}(1)$ *queries*[†]*, and we can find e obliviously with* $\log n/k_{\mathrm{inv}} + 4k_{\mathrm{inv}} + \mathcal{O}(1)$ *queries.*

In general, we cannot hope to obtain results of similar quality for the smaller parameter k_{ainv}. In fact, already for $k_{\mathrm{ainv}} = 1$ any search algorithm needs to query all n elements.

Proposition 2. *For $k_{\mathrm{ainv}} \ge 1$, no algorithm can find e with less than n queries.*

Proof. Consider the family of arrays that are obtained from $[1, \dots, n]$ by moving n to an arbitrary position (possibly leaving it in place); all arrays of this form have $k_{\mathrm{ainv}} \le 1$ since the only possible adjacent inversion is between n and the succeeding element). An adversary may use this family to force any search algorithm to query all n positions when searching for $e = n$: The adversary will answer the first $n-1$ queries by $<$, maintaining that all arrays where n is in any unqueried position are consistent with the answers given so far. (This is easy to see since all other elements are smaller than $e = n$.) □

Fortunately, we can do much better if the target e is guaranteed to be in the correct position relative to sorted order, i.e., if $\mathrm{pos}(e) = \mathrm{rank}(e)$. Note that this restriction still allows us to prove a lower bound on the necessary number of queries that is much larger than all preceding results. We will complement this lower bound by a search algorithm that matches it tightly (up to lower-order terms). Both upper and lower bound hinge on the question of how efficiently (in terms of queries) an algorithm can find a good estimate of $\mathrm{rank}(e)$ by querying the array.

Theorem 9. *Every search algorithm needs at least* $2\sqrt{2nk_{\mathrm{ainv}}} - o(\sqrt{nk_{\mathrm{ainv}}})$ *queries, even if* $\mathrm{pos}(e) = \mathrm{rank}(e)$.[†]

Proof. We describe an adversary that will force any search algorithm to use at least the claimed number of queries, i.e., at least $2\sqrt{2nk_{\text{ainv}}}$ minus lower order terms. The adversary will not fix the actual contents of the array beforehand but will guarantee throughout that there exists an n-element array with at most k_{ainv} adjacent inversions that is consistent with all queries answered so far. We will consider positions of the underlying array to be numbered $1, \ldots, n$, and assume that n is even for convenience.

At high level, the adversary aims to place the target e as close to the middle of the array as possible. Accordingly, his *standard response* to queries to the first half is $<$ whereas it is $>$ for the second half. Eventually, he will be forced to pick a concrete array and position of e that is consistent with all queries. If e is placed in position x in the first half then previous queries between x and the middle have identified elements that are smaller than the target, and that are now found to the right of it. To adhere to the restriction that $\text{pos}(e) = \text{rank}(e)$, the adversary needs to choose an array that has the same number of larger elements to the left of e. He needs to place blocks of such elements between positions left of x that are already queried (and contain smaller elements) without causing too many adjacent inversions; these blocks are called *hidden blocks* in reference to the fact that none of their positions has been queried before. (All of this is symmetric for placement in the second half.)

Let us now give a detailed description of the adversary's strategy. The standard response is $<$ for positions $1, \ldots, \frac{n}{2}$ and $>$ for positions $\frac{n}{2}+1, \ldots, n$. The adversary keeps track of the following values: $p \geq 0$ is the smallest value such that $\frac{n}{2} - p$ has not been queried yet; $q \geq 0$ is the smallest value such that $\frac{n}{2} + 1 + q$ has not been queried yet; ℓ is the number of queries that have been made to positions $1, \ldots, \frac{n}{2} - p - 1$; and r is the number of queries that have been made to positions $\frac{n}{2} + q + 2, \ldots, n-1$. Initially we have $p = q = \ell = r = 0$. Observe that p and q will never decrease, but ℓ and r may decrease upon increases of p or q, respectively. Note that $p + \ell$ and $q + r$ never decrease: E.g., $p + \ell$ counts the queries on $1, \ldots, \frac{n}{2} - p - 1$ and those on $\frac{n}{2} - p + 1, \ldots, \frac{n}{2}$; there is never any previous query for position $\frac{n}{2} - p$ by choice of p. Thus, we also see that $p + q + \ell + r$ is always equal to the total number of queries made so far by the search algorithm.

Concretely, the adversary plans to put the target e either in position $\frac{n}{2} - p$ in the first half or in position $\frac{n}{2} + 1 + q$ in the second half of the array. Queries to the first, respectively second, half of the array may force him to abandon the corresponding option. Abandoning the first, say not putting the target in the first half, allows him to continue answering $<$ in the first half, and to answer $>$ in the second half until he needs to commit to a position $\frac{n}{2} + 1 + q$ in the second half. In other words, the standard response can be continued. Once a standard response would force to abandon the second option he instead needs to commit to an instantiation that puts the target in the previously still feasible half. At this point, no further queries will be guaranteed, since the current query could in principle be to the position of e in the chosen instantiation. Nevertheless, we show that the number $p + q + \ell + r$ of queries of the search algorithm invested will be sufficiently large at this point. To do so, we give a lower bound on $p + \ell$ for the point when the first half is no longer feasible, and an analog lower bound on $q + r$ for the second half.

Intuitively, the adversary aims to have the target as close as possible to the center of the array. Of course, queries to positions around $\frac{n}{2}$ will eventually increase the values of p and q, which are bounds for how close the target e can be to the center. Since the adversary needs to fulfill $\text{pos}(e) = \text{rank}(e)$, the position of e forces the adversary to choose an array with the correct numbers of smaller and larger elements. In particular, if the adversary chooses for example to place the target e in the current position p, then all elements in positions $p+1, \ldots, \frac{n}{2}$ are smaller than the target (due to queries that were already answered); these are exactly p elements. Thus, to balance out the numbers the adversary needs to choose an array such that at least p elements

larger than the target are in positions $1, \ldots, p-1$ (we will go for exactly p and have no additional smaller elements succeeding the target). This is hindered by the fact that ℓ queries were already made on this part, and the fact that every maximal block of larger elements in the first half necessarily ends with an adjacent inversion (either with a smaller element or with the target). This will eventually force the adversary to "give up" on one half of the array or (when this happens the second time) to pick a concrete instantiation with the target placed in the second half.

Now, assume that some query is made by the search algorithm. We will discuss in detail what happens for a query to a position in $1, \ldots, \frac{n}{2}$; queries to $\frac{n}{2}+1, \ldots, n$ are treated symmetrically. For a query to a position in $1, \ldots, \frac{n}{2}$ the adversary tests whether he can answer with $<$ and still maintain existence of a consistent instantiation that places e in the first half. To this end, pretend that the query is answered $<$, update p and ℓ, and check whether there is a consistent instantiation with target in position $\frac{n}{2}-p$ that has at most k_{ainv} adjacent inversions. Note that the adversary does not need to do this optimally. It suffices to have a strategy that causes the claimed number of queries, and he may give up even though there could still be a consistent instantiation.

From the perspective of a sorted array, placing the target in position $\frac{n}{2}-p \in \{1, \ldots, \frac{n}{2}\}$ only conflicts with the queries to positions $\frac{n}{2}-p+1, \ldots, \frac{n}{2}$ that were already answered with $<$, of which there are exactly p. To balance out the total numbers of larger and smaller elements, the adversary checks for an instantiation that puts k blocks of larger elements into positions $1, \ldots, \frac{n}{2}-p-1$, surrounded by smaller elements. This causes k adjacent inversions between the last element of a block and the subsequent element. Moreover, one adjacent inversion exists between the target in position $\frac{n}{2}-p$ and its successor, which must be a smaller element due to queries, barring the trivial case of $p=0$ where a fully sorted array without adjacent inversions suffices. Thus, the adversary may choose $k := k_{\text{ainv}}-1$. In the interest of recycling the argument later, we will do the analysis in terms of k and only plug in $k_{\text{ainv}}-1$ at the very end.

Thus, it suffices to check whether there are k non-overlapping blocks in positions $1, \ldots, \frac{n}{2}-p-1$ of elements that have not been queried yet, whose total size is at least p; placing p larger elements in such blocks will be consistent with queries answered so far. If such a choice of blocks exist then a valid instantiation would be to place p larger elements in these blocks, surrounded by smaller ones, followed by the target in position $\frac{n}{2}-p$, followed by p smaller elements, and finally by $\frac{n}{2}$ larger elements. Observe that the adversary can choose the smaller respectively larger elements freely and, hence, there are no adjacent inversions inside blocks of smaller respectively larger elements. Thus, if the p larger elements that are required can be placed in k blocks in positions $1, \ldots, \frac{n}{2}-p-1$ then there is an instantiation with at most $k+1$ adjacent inversions (k from the blocks and one between the target and its successor). With $k \leq k_{\text{ainv}}-1$ this suffices to maintain the invariant of an existing consistent instantiation.

Let us now derive a lower bound for $p+\ell$ for the case that the adversary cannot find k non-overlapping blocks of unqueried elements in $0, \ldots, \frac{n}{2}-p-1$. Clearly, the total size of *any* k maximal unqueried blocks in $0, \ldots, \frac{n}{2}-p-1$ must be less than p. We can relate n, p, and k by counting the number F of unqueried elements in $1, \ldots, \frac{n}{2}-p-1$ in two different ways. Clearly, because ℓ is the number of queries made to $1, \ldots, \frac{n}{2}-p-1$ we have

$$F = \frac{n}{2} - p - 1 - \ell. \tag{1}$$

On the other hand, letting $B_1, \ldots, B_s$ denote the maximal unqueried blocks in $1, \ldots, \frac{n}{2}-p-1$,

we have

$$F = \sum_{i=1}^{s} |B_i|.$$

For convenience, we explicitly include empty blocks between adjacent queried positions, before a position 1 if it is queried, and after position $\frac{n}{2} - p - 1$ if it is queried; in this way, the number of maximal unqueried blocks is exactly $\ell + 1$. Since any k of these blocks have total size less than p we get $|B_1| + \ldots + |B_k| < p$, $|B_2| + \ldots + |B_{k+1}| < p$, and so on (wrapping around indices larger than s). Summing up these $s = \ell + 1$ inequalities we get

$$(\ell + 1) \cdot p = s \cdot p > k \cdot \sum_{i=1}^{s} |B_i| = k \cdot F.$$

Together with (1) this yields

$$(\ell + 1) \cdot p > k \cdot F = k \cdot \left(\frac{n}{2} - p - 1 - \ell \right).$$

We bring this inequality into a more convenient form to derive a lower bound for $\ell + p$:

$$\begin{aligned} & (\ell + 1) \cdot p > k \cdot \left(\frac{n}{2} - p - 1 - \ell \right) \\ \Leftrightarrow \quad & (\ell + 1) \cdot p + k \cdot (p + \ell) > k \cdot \left(\frac{n}{2} - 1 \right) \end{aligned} \tag{2}$$

The left hand side is upper bounded by

$$\left(\frac{p + \ell + 1}{2} \right)^2 + k \cdot (p + \ell) \geq (\ell + 1) \cdot p + k \cdot (p + \ell), \tag{3}$$

since, for any $x, y \geq 0$ we have

$$\begin{aligned} & \left(\frac{x - y}{2} \right)^2 \geq 0 \\ \Rightarrow \quad & \frac{x^2}{4} - \frac{xy}{2} + \frac{y^2}{4} \geq 0 \\ \Rightarrow \quad & \frac{x^2}{4} + \frac{xy}{2} + \frac{y^2}{4} \geq xy \\ \Rightarrow \quad & \left(\frac{x + y}{2} \right)^2 \geq xy; \end{aligned}$$

we use it with $x = \ell + 1$ and $y = p$ to obtain (3).

Combining (2) with (3) yields

$$\begin{aligned} & \left(\frac{p + \ell + 1}{2} \right)^2 + k \cdot (p + \ell) > k \cdot \left(\frac{n}{2} - 1 \right), \\ \Leftrightarrow \quad & \left(\frac{p + \ell + 1}{2} \right)^2 + k \cdot (p + \ell + 1) - k - k \cdot \left(\frac{n}{2} - 1 \right) > 0, \\ \Leftrightarrow \quad & \left(\frac{p + \ell + 1}{2} \right)^2 + k \cdot (p + \ell + 1) - k \cdot \frac{n}{2} > 0, \qquad (4) \\ \Leftrightarrow \quad & (p + \ell + 1)^2 + 4k \cdot (p + \ell + 1) - 2kn > 0. \qquad (5) \end{aligned}$$

Using $p+\ell+1>1$ and computing the roots of (5) wrt. $p+\ell+1$ we arrive at

$$p+\ell+1 > -2k+\sqrt{4k^2+2kn}.$$
$$\Rightarrow \quad p+\ell \geq \sqrt{4k^2+2kn}-2k.$$

Thus, if the proposed instantiation is not possible using k blocks of larger elements then we have

$$p+\ell \geq \sqrt{4k^2+2kn}-2k = \sqrt{2kn}-o(\sqrt{kn}).$$

Similarly, if there is no feasible instantiation placing the target at $\frac{n}{2}+1+q$ in the second half of the array then we can prove that $q+r \geq \sqrt{4k^2+2kn}-2k = \sqrt{2kn}-o(\sqrt{kn})$. We give the calculations here for completeness.

Let us check first that the adversary can use the same number $k = k_{\text{ainv}}-1$ of hidden blocks: His goal is let the second half, i.e., positions $\frac{n}{2}+1,\ldots,n$, contain (in order) q larger elements, the target in position $\frac{n}{2}+1+q$, and larger elements interspersed by up to k blocks of smaller elements. As before, elements within the group of larger/smaller elements can be assumed to be sorted, thus, adjacent inversions are only possible before the target (if preceded by a larger element), or before a smaller element (if preceded by the target or an element larger than the target). Clearly, we get an adjacent inversion at each of the k blocks, a single inversion between the q larger elements and the target, and no further adjacent inversions. Hence, $k = k_{\text{ainv}}-1$ hidden blocks are a feasible choice.

If placing at $\frac{n}{2}+1+q$ is infeasible then in particular there are no k blocks of unqueried elements in $\frac{n}{2}+q+2,\ldots,n$ of total size at least q. Again, using that the number F' of unqueried elements in $\frac{n}{2}+q+2,\ldots,n$ is equal to $\frac{n}{2}-q-1-r$, but also equal to the total size of the $s'=r+1$ maximal unqueried blocks $B'_1,\ldots,B'_{s'}$ (including size-zero blocks), we get

$$(r+1)\cdot q = s'\cdot q > k\cdot\sum_{i=1}^{s'}|B'_i| = k\cdot F' = k\cdot\left(\frac{n}{2}-q-1-r\right).$$

This implies

$$(r+1)\cdot q + k\cdot(q+r) > k\cdot\left(\frac{n}{2}-1\right).$$

At this point it is obvious that we arrive at the same lower bound for $q+r$ as we had for $p+\ell$, i.e., as claimed above.

Thus, we conclude that if the adversary gets to use instantiations with k hidden blocks (as above) then he can enforce a total of at least

$$p+\ell+q+r \geq 2\sqrt{4k^2+2kn}-4k = 2\sqrt{2kn}-o(\sqrt{kn})$$

queries. For the case of $k = k_{\text{ainv}}-1$ we get a lower bound of $2\sqrt{2k_{\text{ainv}}n}-o(\sqrt{k_{\text{ainv}}n})$ as claimed. □

We now describe an algorithm that achieves the optimal number of queries (up to lower order terms). Note that the algorithm requires knowledge of k_{ainv}.

Theorem 10. *We can find e using $2\sqrt{2nk_{\text{ainv}}}+o(\sqrt{nk_{\text{ainv}}})$ queries if $\text{pos}(e)=\text{rank}(e)$.*

Proof. For the description of our algorithm it will be convenient to take the array as having positions numbered $0,\ldots,n-1$. As before, a query for some position i will yield $<$, $>$, or $=$ depending on whether $A[i]<e$, $A[i]>e$, or $A[i]=e$. The assumption that $\text{pos}(e)=\text{rank}(e)$

will be crucial; the algorithm will attempt to get a good estimate for rank(e) and then query a certain range around the estimated position.

The algorithm first fixes a block size of the form $p = c \cdot \sqrt{\frac{n}{k_{\text{ainv}}}}$ for some constant c that we will fix later, and then queries positions $0, p+1, 2(p+1), \ldots, \lceil \frac{n-1}{p} \rceil (p+1), n-1$. We refer to these positions as the *grid*. In this way, unqueried blocks of size (at most) p remain, with every block sandwiched between two grid positions. We will refer to these blocks according to query outcomes to the adjacent grid positions: $<>$-blocks, $<<$-blocks, $><$-blocks, and $>>$-blocks. We use $\sharp(xy)$ to denote the number of xy-blocks for $x, y \in \{<, >\}$. We assume that the target e is not at a grid position, since otherwise the claimed bound holds trivially.

Intuitively, since there are at most k_{ainv} adjacent inversions, there can only be a limited number of $>>$-blocks containing elements smaller than e and of $<<$-blocks containing larger elements: Either constellation leads to at least one adjacent inversion inside the block. Moreover, every $><$-block must contain an adjacent inversion, which upper bounds their number by k_{ainv} (we will give a better bound later). The number of $<>$-blocks is at most equal to the number of $><$-blocks plus one, because there must be an $><$-block somewhere between any two $<>$-blocks.

The algorithm now tries to estimate the position of e in the array. Crucially, because pos(e) = rank(e), there must be exactly pos(e) elements smaller than e in the array. We denote by $q(x)$ the number of queries to grid positions that returned x, for $x \in \{<, >, =\}$. We further denote by $\eta(xy)$, for $x, y \in \{<, >\}$, the number of adjacent inversions involving an element of a xy-blocks, including adjacent inversions between an element of the block and an adjacent grid position. We denote by $N(xy)$, for $x, y \in \{<, >\}$, the total number of positions in xy-blocks (not counting the queries adjacent to the block).

Observe that pos(e) is equal to the number of elements smaller than e in the array. We refer to these elements simply as *small* elements, as opposed to *large* elements that are larger than e. We want to bound the number of small elements in order to obtain a range of positions that our algorithm needs to search. By definition, we have $q(<)$ small elements in grid positions. We now give upper and lower bounds for the number of small elements in each type of block in terms of $\sharp(xy)$ and $\eta(xy)$. It can be observed that upper and lower bound are attained by blocks not containing e, hence we will tacitly ignore this case.

$<<$-blocks: The maximum number $N(<<)$ of smaller elements in $<<$-blocks is attained if there are no adjacent inversions. If there is at least one adjacent inversion in a $<<$-block, it can have anywhere between 0 and block size (at most p) small elements: If the large elements form a single block then there is exactly one adjacent inversion between the last element of the block and the succeeding element (possibly a grid position). Overall, the number of small elements in $<<$-blocks is between $N(<<) - \eta(<<) \cdot p$ and $N(<<)$.

$>>$-blocks: The minimum number of small elements in $>>$-blocks is attained if there are no adjacent inversions; in this case there are no smaller elements in these blocks. A $>>$-block with at least one adjacent inversion can contain any number between 0 and the block size (at most p) of small elements: A consecutive block of small elements causes a single adjacent inversion at its start. Overall, the number of small elements in $>>$-blocks is between 0 and $\eta(>>) \cdot p$.

$><$-blocks: Every $><$-block contains at least one adjacent inversion. In such a block, even without any further adjacent inversions, we can have between 0 and the block size (at most p) smaller elements: The block may contain all large elements followed by all small

ones, and have a single adjacent inversion between the last large and first small element. Overall, the number of small elements in $><$-blocks is between 0 and $\sharp(><) \cdot p$. Since there is at least one adjacent inversion per $><$-block we have $\sharp(><) \leq \eta(><)$ and the upper bound becomes $\eta(><) \cdot p$.

$<>$-blocks: In $<>$-blocks we can have any number of small elements followed by large ones (with the total being at most the block size p) without any adjacent inversions. We noted already that the number of these blocks is at most $\sharp(><) + 1$. Overall, the number of small elements in $<>$-blocks is between 0 and $\sharp(<>) \cdot p \leq (\sharp(><)+1) \cdot p \leq (\eta(><)+1) \cdot p$.

In total, we get that there are at least

$$q(<) + N(<<) - \eta(<<) \cdot p$$

and at most

$$q(<) + N(<<) + \eta(>>) \cdot p + \eta(><) \cdot p + (\eta(><) + 1) \cdot p$$

elements that are smaller than e.

Since $\mathrm{pos}(e)$ is equal to the number of small elements in the array, the gap (plus one) between these two bounds is the number of positions that the algorithm has to query in order to find e or be sure that it is not present. The gap (difference) is upper-bounded by

$$\begin{aligned}
&(q(<) + N(<<) + \eta(>>) \cdot p + \eta(><) \cdot p + (\eta(><) + 1) \cdot p) \\
&\quad - (q(<) + N(<<) - \eta(<<) \cdot p) + 1 \\
&= \eta(>>) \cdot p + \eta(><) \cdot p + (\eta(><) + 1) \cdot p + \eta(<<) \cdot p + 1 \\
&= \eta(>>) \cdot p + \eta(<<) \cdot p + 2\eta(><) \cdot p + p + 1.
\end{aligned}$$

Since $\eta(>>) + \eta(<<) + \eta(><) + \eta(<>) \leq k_{\mathrm{ainv}}$ and since all values are non-negative, this expression is maximized for $\eta(><) = k_{\mathrm{ainv}}$ and $\eta(<>) = \eta(<<) = \eta(>>) = 0$, i.e., if there are no adjacent inversions inside $<<$-blocks, $>>$-blocks, and $<>$-blocks. We get a range of at most $(2k_{\mathrm{ainv}} + 1) \cdot p + 1$ positions to search, which coincides with the claim of the theorem, since $p = c\sqrt{n/k_{\mathrm{ainv}}}$. Unfortunately, while the algorithm knows k_{ainv} and can thus compute the *number of elements* in the range it has to search, it does not know exactly *where the range starts or ends* because it does not have access to the values of η for each block and cannot compute the value of the upper or lower bound.

Since the search range is maximized when all inversions fall in $><$-blocks it makes sense to refine our initial grid in order to get rid of $><$-blocks altogether. We can do this by running a binary search on each $><$-block to find an adjacent inversion in at most $1 + \log p$ queries. To do this, we simply query the center element and recurse on the left subblock if it is small and on the right subblock if it is large, until we are left with a subblock containing a large element followed by a small one. By extending our initial grid by the additional query positions, we replace each $><$-block by some number of $>>$-blocks, an empty $><$-block (the adjacent inversion), and some number of $<<$-blocks. We update the values of $\sharp$, N, η and q accordingly.

Overall, our algorithm spends at most $\sharp(><) \cdot (1 + \log p)$ queries on the refinement of the grid, resulting in all $><$-blocks being empty. Thus, we get tighter bounds for the candidate range for $\mathrm{pos}(e)$ by setting the block size of $><$-blocks to 0 (instead of p), which eliminates the term $\eta(><) \cdot p$. Note that while we introduced additional $<<$-blocks and $>>$-blocks and may thus have increased $\eta(<<)$ and $\eta(>>)$, the additional blocks contain fewer than p elements. Nevertheless, we can use the generous bound of p for the size of all blocks, since the block size

only appears negatively in the lower bound and positively in the upper bound. We retain the lower bound of

$$q(<) + N(<<) - \eta(<<) \cdot p$$

and get a new upper bound of

$$q(<) + N(<<) + \eta(>>) \cdot p + (\eta(><) + 1) \cdot p,$$

Using that $-k_{\text{ainv}} \leq -\eta(<<) \leq 0$ and $0 \leq \eta(>>) + \eta(><) \leq k_{\text{ainv}}$ we conclude that the number of smaller elements is lower bounded by

$$q(<) + N(<<) - k_{\text{ainv}} \cdot p$$

and upper bounded by

$$q(<) + N(<<) + k_{\text{ainv}} \cdot p + p.$$

Since the bounds depend only on values that are known to the algorithm, it can simply query all positions in this range. Since $\text{pos}(e) = \text{rank}(e)$, if the target e is contained in the array then it must be in the position that equals the number of smaller elements. Thus, it suffices to query the $2k_{\text{ainv}}p + p + 1$ elements between the above bounds. Note that small savings are possible here because some of the positions have been queried previously, but we will not analyze this.

Thus, the overall number of queries needed to establish the initial grid setup, for its refinement, and for the final sweep of the candidate range for $\text{pos}(e)$ is at most

$$\frac{n}{p+1} + 2 + \sharp(><) \cdot (1 + \log p) + 2k_{\text{ainv}} \cdot p + p + 1 = \frac{n}{p+1} + 2k_{\text{ainv}} \cdot p + o(k_{\text{ainv}}p), \qquad (6)$$

where $\frac{n}{p+1} + 2$ upper bounds the number queries needed to establish the initial grid that partitions the array into unqueried blocks of size at most p each. Rounding up slightly, we are left with choosing c in $p = c \cdot \sqrt{\frac{n}{k_{\text{ainv}}}}$ in order to minimize

$$\frac{n}{p+1} + 2k_{\text{ainv}} \cdot p \leq \frac{n}{p} + 2k_{\text{ainv}} \cdot p. \qquad (7)$$

In other words, after plugging in $p = c \cdot \sqrt{\frac{n}{k_{\text{ainv}}}}$, we seek c that minimizes

$$\frac{n}{c \cdot \sqrt{\frac{n}{k_{\text{ainv}}}}} + 2k_{\text{ainv}} \cdot c \cdot \sqrt{\frac{n}{k_{\text{ainv}}}} = \frac{1}{c}\sqrt{nk_{\text{ainv}}} + 2c \cdot \sqrt{nk_{\text{ainv}}} = \left(\frac{1}{c} + 2c\right) \cdot \sqrt{nk_{\text{ainv}}}.$$

Choosing $c = \frac{1}{\sqrt{2}}$ yields $(\frac{1}{c} + 2c) = 2\sqrt{2}$ and the claimed bound of $2\sqrt{2nk_{\text{ainv}}} + o(\sqrt{nk_{\text{ainv}}})$ queries for finding the target e. □

4.3. Edit distances

In this section, we consider parameters that bound the number of elementary array modifications needed to sort the given array A. More precisely, a *replacement* is the operation of replacing one element with a new element, and we let k_{rep} be the (minimum) number of replacements needed to obtain a sorted array. A *swap* is the exchange of the content of two array positions, and k_{swap} is the number of swaps needed to sort A. We let k_{aswap} be the number of swaps of pairs of neighboring elements needed to sort A. A *move* is the operation of removing an element

and re-inserting it after a given position i, shifting all elements between old and new position by one. We let k_{mov} be the number of moves needed to sort A.

Clearly, starting from a sorted array, we can move e to any position, without using more than a single move or swap, or two replacements involving e. To find e we then have to query the entire array.

Proposition 3. *For $k_{\text{mov}} \geq 1$, $k_{\text{swap}} \geq 1$, or $k_{\text{rep}} \geq 2$, no algorithm can find e with less than n queries in general.*

We can obtain significantly improved bounds if the element e remains at its correct position relative to the sorted array. Recall that we can interpret k_{faults} as a measure of disorder via $k_{\text{faults}}(e) = |i : (i < \text{pos}(e) \wedge A[i] > e) \vee (i > \text{pos}(e) \wedge A[i] < e)|$.

Lemma 2. *If $\text{rank}(e) = \text{pos}(e)$, then $k_{\text{faults}}(e) \leq \min\{2k_{\text{rep}}, 4k_{\text{swap}}, 2k_{\text{mov}}\}$.*

Proof. Consider an array A with $k_{\text{rep}} = k$. We get a modified array A' with $k_{\text{rep}} = k'$ by switching out all elements smaller than e in A with a common element $e_< < e$ and all larger elements by $e_> > e$. Let $r_j = (i_j, e_j), j \in \{1, \ldots, k\}$ be k replacements that transform A into a sorted array. We define $r'_j = (i_j, 0)$ if $e_j < e$ and $r'_j = (i_j, 2e)$ otherwise. Clearly, the replacements $r'_j = (i_j, e_j), j \in \{1, \ldots, k\}$ transform A' into a sorted array, and, hence, $k' \leq k$. Let m be the number of entries left of e that contain $2e$. Since $\text{rank}(e) = \text{pos}(e)$, we have that m is also the number of entries equal 0 right of e. It is clear that $k' \geq m$, since we have m disjoint pairs of elements in the wrong relative order that need to be repaired by replacing at least one of the two. On the other hand, in both A and A', we have $k_{\text{faults}}(e) = 2m \leq 2k' \leq 2k$.

Now assume A has $k_{\text{swap}} = k$ and let $s_j = (i_j, i'_j), j \in \{1, \ldots k\}$ be the k swaps (given by the indices of the swapped elements) that transform A into a sorted array. Clearly, the same sequence of swaps turns A' into a sorted array, and, hence, $k' \leq k$, where A' has $k_{\text{swap}} = k'$. As before, we have m disjoint pairs of elements in the wrong relative order that need to be repaired by switching at least one of the two elements. Each switch can repair two such pairs, and we thus have $k' \geq m/2$. In both A and A', we have $k_{\text{faults}}(e) = 2m \leq 4k' \leq 4k$.

Similarly, the transformation from A to A' does not increase k_{mov}. Again, we have m disjoint pairs of elements in the wrong relative order that need to be repaired by moving at least one of the two. Hence, $k_{\text{faults}}(e) = 2m \leq 2k_{\text{mov}}$. □

With this lemma, we can translate the upper bounds of any algorithm for k_{lies}. Before we do, we introduce another measure of disorder, that turns out to be closely related to k_{mov}. We define the parameter k_{seq} to be such that $n - k_{\text{seq}}$ is the length of a longest nondecreasing subsequence in A. It turns out that $k_{\text{mov}} = k_{\text{seq}}$ (cf. Proposition 5), and we can thus include this parameter in our upper bound.

Theorem 11. *Let $f\colon \mathbb{N}^2 \to \mathbb{N}$. If $\text{rank}(e) = \text{pos}(e)$, and we can find e with $f(n, k_{\text{lies}})$ queries, then we can find e obliviously with $\min\{f(n, 2k_{\text{rep}}), f(n, 4k_{\text{swap}}), f(n, 2k_{\text{mov}}), f(n, 2k_{\text{seq}})\}$ queries.*

We can also carry over the lower bound from Corollary 1.

Corollary 4. *For every $c \in \mathbb{N}$ and $\text{pos}(e) = \text{rank}(e)$, no algorithm operating on the search tree can find e with less than $\log n + ck$ queries in general, for $k \in \{k_{\text{rep}}, k_{\text{swap}}, k_{\text{mov}}, k_{\text{seq}}\}$.*[†]

Proof. Because of $\text{pos}(e) = \text{rank}(e)$, we group all indices in A with at least one wrong query into m disjoint pairs with one element left of e and one element right of e in each pair.

Starting from a sorted array, we can clearly produce the array A by swapping each pair. This requires $k_{\text{faults}}/2$ swaps, or k_{faults} element moves or replacements. We get $k_{\text{faults}} \geq \max\{k_{\text{rep}}, k_{\text{swap}}, k_{\text{mov}}\}$. Corollary 1 together with $k_{\text{mov}} = k_{\text{seq}}$ (Proposition 5) thus implies the claim. □

Finally, we immediately obtain bounds for k_{aswap} from Corollary 3, because $k_{\text{aswap}} = k_{\text{inv}}$ (cf. Proposition 6).

Corollary 5. *Every search algorithm needs at least* $\log n/k_{\text{aswap}} + 2k_{\text{aswap}} + \mathcal{O}(1)$ *queries*†*, and we can find e obliviously with* $\log n/k_{\text{aswap}} + 4k_{\text{aswap}} + \mathcal{O}(1)$ *queries.*

4.4. Block edit distances

In this section we consider the parameters k_{bswap}, k_{rbswap}, and k_{bmov}, which bound the number of block edit operations needed to sort A. A *block* is defined to be a subarray $A[i, i+1, \ldots, j]$ of consecutive elements. A *block swap* is the operation of exchanging a subarray $A[i, \ldots, j]$ with a subarray $A[i', \ldots, j']$ and vice versa, where $i < j < i' < j'$. Note that a block swap may affect the positions of other elements in case that the two blocks are of different sizes. The parameter k_{bswap} bounds the number of block swaps needed to sort A. For k_{rbswap} we only allow block swaps restricted to pairs of blocks of equal sizes. Finally, for k_{bmov} one of the two blocks must be empty, i.e., only block moves are allowed. For all three parameters one can easily prove that, without further restrictions, search algorithms need to query all positions of an array to find the target.

Proposition 4. *For $k_{\text{bswap}} \geq 1$, $k_{\text{rbswap}} \geq 1$, or $k_{\text{bmov}} \geq 1$, no algorithm can find e with less than n queries.*

Proof. For k_{bmov} and k_{bswap} consider the family of arrays obtained from $[1, \ldots, n]$ by moving n to an arbitrary position; arrays of this form have $k_{\text{bmov}} = k_{\text{bswap}} \leq 1$. An adversary may answer the first $n-1$ queries by $<$ while maintaining that placing $e = n$ in any of the unqueried positions is consistent with all given answers.

For k_{rbswap} consider the family of arrays obtained from $[1, \ldots, n]$ by swapping n with any element (or keeping it in place; such arrays have $k_{\text{rbswap}} \leq 1$. An adversary may again answer the first $n-1$ queries by $<$ while maintaining that placing $e = n$ in any unqueried position is consistent with the given answers. □

Complementing this lower bound, for all of three parameters, an upper bound of $\mathcal{O}(\sqrt{nk})$ for finding e when $\text{pos}(e) = \text{rank}(e)$ follows immediately from the results for k_{ainv} of Section 4.2 and the fact that $k_{\text{ainv}} \leq 2k_{\text{bswap}}$ (Proposition 16) and $k_{\text{bswap}} \leq \min\{k_{\text{rbswap}}, k_{\text{bmov}}\}$ (Propositions 9 and 11). By inspecting the upper and lower bounds proved for k_{ainv}, and adapting the proofs, we are able to obtain tight leading constants in the upper and lower bounds for k_{bmov} and k_{bswap}, and leading constants within a factor of $\sqrt{2}$ for k_{rbswap}. First, we adapt the lower bound for k_{ainv} (Theorem 9) to k_{rbswap} and k_{bswap}.

Theorem 12. *Every search algorithm needs at least* $2\sqrt{2nk_{\text{rbswap}}} - o(\sqrt{nk_{\text{rbswap}}})$ *queries to find e, even if* $\text{pos}(e) = \text{rank}(e)$.†

Proof. We use the same adversary setup as in Theorem 9 but we need to now make sure that the adversary maintains existence of a suitable instantiation that is at most k_{rbswap} restricted block swaps away from being sorted. We will show that the same strategy can be used, but with $k = k_{\text{rbswap}}$ hidden blocks.

We only discuss the case that the adversary commits to an instantiation with the target e in the first half of the array and having k blocks of large elements between already queried positions. By the analysis from the proof of Theorem 9 we know that when the adversary commits to putting the target into position $\frac{n}{2}-p$ in the first half, there exist k non-overlapping blocks of unqueried elements in the first half, and with total size at least p. The instantiation consists of (in this order) small elements in positions $1,\ldots,\frac{n}{2}-p-1$, interspersed with k blocks of large elements of total length p; the target e in position $\frac{n}{2}-p$; p small elements in positions $\frac{n}{2}-p+1,\ldots,\frac{n}{2}$; and $\frac{n}{2}$ large elements in positions $\frac{n}{2}+1,\ldots,n$. Clearly, this can be turned by $k=k_{\text{rbswap}}$ restricted block swaps into an array with small elements, followed by the target, and followed by large elements. Since the adversary does not need to report the numerical values, the actual numbers can be chosen with the instantiation in such a way that the latter interval is sorted.

The lower bound is thus obtained by plugging in $k=k_{\text{rbswap}}$ into the lower bound of

$$2\sqrt{4k^2+2kn}-4k = 2\sqrt{2kn}-o(\sqrt{kn}).$$

For $k=k_{\text{rbswap}}$ this yields the claimed lower bound. □

A lower bound of $2\sqrt{2nk_{\text{bswap}}}-o(\sqrt{nk_{\text{bswap}}})$ now follows from $k_{\text{bswap}}\le k_{\text{rbswap}}$ (Proposition 9), but a better lower bound is obtained in the following theorem.

Theorem 13. *Every search algorithm needs at least* $4\sqrt{nk_{\text{bswap}}}-o(\sqrt{nk_{\text{bswap}}})$ *queries to find* e *even if* $\text{pos}(e)=\text{rank}(e)$.[†]

Proof sketch. Again we use the same adversary setup as in Theorem 9 and this time focus on (unrestricted) block swaps rather than restricted block swaps. To get a stronger lower bound, the adversary will use more hidden blocks, namely $2k_{\text{bswap}}-1$ hidden blocks. This of course requires using less than one block swap for each hidden block (unlike the previous case of restricted block swaps). The following claim shows an appropriate construction of k hidden blocks with only $\frac{k+1}{2}$ block swaps. It is formulated for the case of placing e in the second half of an array but the other case is symmetric. Positions are numbered 1 through n.

Claim 1. *Let* $n\in\mathbb{N}$*, and let* $\alpha_1,\ldots,\alpha_k,\beta_0,\beta_1,\ldots,\beta_k\in\mathbb{N}$ *with* $\sum\alpha_i=q$ *and* $\sum\beta_i=\frac{n}{2}-q-1$*. There is an array* A *with* $k_{\text{bswap}}(A)\le\lceil\frac{k+1}{2}\rceil$ *that contains (in order) the following elements:* $\frac{n}{2}$ *elements smaller than* e*,* q *elements larger than* e*, element* e*, and an alternating sequence of blocks of larger and smaller elements of sizes* $\beta_0,\alpha_1,\beta_1,\ldots,\alpha_k,\beta_k$ *(with* β_i *the sizes of blocks of larger elements).*

Proof. Start with any sorted array A' that contains element e in position $\frac{n}{2}+q+1$. Accordingly, with positions numbered 1 to n, array A' has exactly $\frac{n}{2}+q=\frac{n}{2}+\sum\alpha_i$ elements that are smaller than e and exactly $\frac{n}{2}-q-1=\sum\beta_i$ elements that are larger than e. We will construct the desired array A from A' by a sequence of at most $\lceil\frac{k+1}{2}\rceil$ block swaps; we start with $A:=A'$.

As a first block swap, exchange $A[\frac{n}{2},\frac{n}{2}-q-1]$ with $A[\frac{n}{2}+q+\beta_0,\frac{n}{2}+q+\beta_0+q]$. In A we now have the following structure: $\frac{n}{2}$ smaller elements, q larger elements, element e, β_0 larger elements, q smaller elements, and $\beta_1+\ldots+\beta_k$ larger elements. All further operations will only be among these final two groups of elements. For convenience, we discuss the remaining operations on the subarray $\hat{A}$ containing only the final $q=\alpha_1+\ldots+\alpha_k$ smaller elements followed by $\beta_1+\ldots+\beta_k$ larger elements. Clearly, operations turning $\hat{A}$ into an alternating sequence of smaller and larger blocks of sizes $\alpha_1,\beta_1,\alpha_2,\ldots,\alpha_k,\beta_k$ can also be applied to get A into the required form. We show how to do this with $\lceil\frac{k-1}{2}\rceil$ block swaps.

If $k = 1$ then $\hat{A}$ has already the required form and we use $0 = \lceil \frac{k-1}{2} \rceil$ block swaps; getting A into correct form thus used $1 = \lceil \frac{k+1}{2} \rceil$ block swaps. If $k = 2$ then we need to transform $\alpha_1 + \alpha_2$ small elements followed by $\beta_1 + \beta_2$ large ones into pattern $\alpha_1, \beta_1, \alpha_2, \beta_2$, which can be done by swapping the last α_2 small elements with the first β_1 large ones; in total we use two swaps on A.

For $k \geq 3$ we show how one block swap reduces the remaining subproblem to one with $k' = k - 2$. We have $\alpha_1 + \ldots + \alpha_k$ small elements followed by $\beta_1 + \ldots + \beta_k$ large ones, and need to reach pattern $\alpha_1, \beta_1, \alpha_2, \ldots, \alpha_k, \beta_k$. We will swap $\hat{A}[\alpha_1 + 1, \ldots, \alpha_1 + \alpha_k]$ with $\hat{A}[|\hat{A}| - \beta_k - \beta_1, |\hat{A}| - \beta_k - 1]$, i.e., we are swapping α_k small elements with β_1 small ones. The result is that $\hat{A}$ now contains (in order) the following blocks: (1) α_1 small elements (not moved this time), (2) β_1 large elements (just swapped), (3) $\alpha_2 + \ldots + \alpha_{k-1}$ small elements (not swapped this time, but possibly shifted)[3], (4) $\beta_2 + \ldots + \beta_{k-1}$ large elements (not swapped), (5) α_k small elements (just swapped), and (6) β_k large elements (never moved).

Observe that the remaining problem now becomes to transform a subarray $\tilde{A}$ with $\alpha_2 + \ldots + \alpha_{k-1}$ small elements followed by $\beta_2 + \ldots + \beta_{k-1}$ large ones into one with pattern $\alpha_2, \beta_2, \alpha_3, \ldots, \alpha_{k-1}, \beta_{k-1}$. This part is situated in $\hat{A}[\alpha_1 + \beta_1 + 1, |\hat{A}| - \alpha_k - \beta_k - 1]$ and hence also in A, and performing the block swaps on this part does not affect the already correctly placed elements. Thus, overall we need at most $\lceil \frac{k+1}{2} \rceil$ block swaps, as claimed. □

Thus, for a lower bound in terms of the number k_{bswap} of block swaps the adversary can use $k = 2(k_{\text{bswap}} - 1)$ hidden blocks: An array instantiation with the k hidden blocks in the required positions costs him only $\lceil \frac{k+1}{2} \rceil = k_{\text{bswap}}$ block swaps. Using this, we can plug $k = 2(k_{\text{bswap}} - 1)$ into the lower bound of

$$2\sqrt{4k^2 + 2kn} - 4k = 2\sqrt{2kn} - o(\sqrt{kn})$$

and obtain the claimed lower bound of $4\sqrt{nk_{\text{bswap}}} - o(\sqrt{nk_{\text{bswap}}})$. □

Now, we directly get a lower bound in terms of the number k_{bmov} of block moves since, using $k_{\text{bmov}} \leq 2k_{\text{bswap}}$, a more efficient search would otherwise violate the lower bound for k_{bswap}.

Corollary 6. *Every search algorithm needs at least* $2\sqrt{2nk_{\text{bmov}}} - o(\sqrt{nk_{\text{bmov}}})$ *queries to find* e *even if* $\text{pos}(e) = \text{rank}(e)$.†

A matching upper bound for k_{bswap}, in the sense of $4\sqrt{nk_{\text{bswap}}} + o(\sqrt{nk_{\text{bswap}}})$, follows immediately from the fact that $k_{\text{ainv}} \leq 2k_{\text{bswap}}$ (Proposition 16). The same bound can be obtained relative to the number of block moves, using $k_{\text{ainv}} \leq 2k_{\text{bmov}}$ or $k_{\text{bswap}} \leq k_{\text{bmov}}$, but a tight upper bound of $2\sqrt{2nk_{\text{bmov}}} + o(\sqrt{nk_{\text{bmov}}})$ is proved in Theorem 14 below. For the number of restricted block swaps, i.e., swapping only blocks of the same size (never incurring any shifts), we get an upper bound of $4\sqrt{nk_{\text{rbswap}}} + o(\sqrt{nk_{\text{rbswap}}})$ from $k_{\text{bswap}} \leq k_{\text{rbswap}}$ (Proposition 9) but this is not tight regarding the leading constant but asymptotically tight.

Corollary 7. *We can find* e *using* $4\sqrt{nk_{\text{bswap}}} + o(\sqrt{nk_{\text{bswap}}})$ *queries if* $\text{pos}(e) = \text{rank}(e)$.

Corollary 8. *We can find* e *using* $4\sqrt{nk_{\text{rbswap}}} + o(\sqrt{nk_{\text{rbswap}}})$ *queries if* $\text{pos}(e) = \text{rank}(e)$.

Theorem 14. *We can find* e *using* $2\sqrt{2nk_{\text{bmov}}} + o(\sqrt{nk_{\text{bmov}}})$ *queries if* $\text{pos}(e) = \text{rank}(e)$.

[3] The shifting of elements will not be relevant here but we mention it once to point out that it is not overlooked.

Proof. The idea for the proof is to revisit the upper bound for number k_{ainv} of adjacent inversions and observe that only the partition into blocks of elements smaller or greater than e matter (in addition to the single element e). Such a partition can have at most $2k_{\text{ainv}} + 2$ blocks because every block of large elements is followed by an adjacent inversion. We can get a similar bound in terms of the number k_{bmov} of block moves, which allows us to conclude the analog bound of $2\sqrt{2nk_{\text{bmov}}} + o(\sqrt{nk_{\text{bmov}}})$ for k_{bmov}. (Note that any sequence of k block moves that sorts the given array, can be reversed into one turning a sorted array into the given one.)

Claim 2. *Starting from a sorted array containing at most a single copy of e, any sequence of at most k block moves gives a partition into at most $2k+2$ maximal blocks of elements smaller, respectively larger than e, and possibly a unit block for the target.*

Proof. There is nothing to prove if there are only small or only large elements (but we give the proof independent of presence of e). For convenience, assume that A only contains elements x, y, and e with $x < e$ and $y > e$ (i.e. possibly many copies of x and y, and at most a single copy of e). We consider a single block move and show that it increases the number of maximal x- or y-blocks by at most 2. Concretely, we show that the number of alternations between blocks increases by at most two (where we also count alternations between x- or y-blocks with e).

Say that we have $A = (\ldots, a, b, \ldots, c, d, \ldots, e, f, \ldots)$ and we move block $(d, \ldots, e)$ between a and b, obtaining $A' = (\ldots, a, d, \ldots, e, b, \ldots, c, f, \ldots)$. Note that $a, b, c, d, e, f \in \{x, y, e\}$. We claim that this block move cannot increase the number of block alternations by three. Assume, for contradiction, that it indeed increases this number by three. Since there are only three new adjacencies, it follows that (a, d), (e, b), and (c, f) must be alternations, i.e., $a \neq d$, $e \neq b$, and $c \neq f$. Similarly, we may not have removed alternations (or else the increase is at most two) so $a = b$, $c = d$, and $e = f$. It follows immediately that none of $a, \ldots, f$ is equal to e since that would imply having at least two copies. Accordingly, $a, b, c, d, e, f \in \{x, y\}$ and we will use $\overline{x} = y$ and $\overline{y} = x$, which allows us to replace, e.g., $a \neq d$ by $a = \overline{d}$. Thus, we get six equalities that together yield $a = \overline{d} = \overline{c} = f = e = \overline{b} = \overline{a}$; a contradiction. It follows that no block move can create more than two additional block alternations, i.e., no block move can increase the number of maximal blocks by more than two.

If e is present then the sorted array has small elements, followed by e, followed by large elements; a total of three blocks. This increases to at most $2k + 3$ blocks after k block moves, one of which is e. If e is not present then we go from 2 blocks to at most $2k+2$. This completes the proof of the claim. □

It follows that arrays that can be sorted with at most k_{bmov} block moves have at most $k_{\text{bmov}} + 1$ pairs consisting of a large element followed by a small element: Such pairs can only occur between different blocks, neither of which is the block containing e. There are at most $2k_{\text{bmov}}+2$ other blocks and hence at most $2k_{\text{bmov}}+1$ alternations between such blocks. Clearly, only every second block alternation can be from larger to smaller element, giving the claimed number of at most $k_{\text{bmov}} + 1$ adjacent alternations between an element larger than e and an element smaller than e. Alternations of this type are the defining quantity for the algorithm given in Theorem 10 for parameter k_{ainv}; there we used that this number is at most k_{ainv} since they are a special case of adjacent alternations. This yields the claimed upper bound of $2\sqrt{2nk_{\text{bmov}}} + o(\sqrt{nk_{\text{bmov}}})$, completing the proof. □

5. Conclusion

We presented upper and lower bounds for the worst-case query complexity of comparison-based search algorithms that are robust to persistent and temporary read errors, or are adaptive to partially disordered input arrays. For many cases we gave algorithms that are optimal up to lower order terms. In addition, many of the algorithms are oblivious to the value of the parameter quantifying errors/disorder, assuming the target element is present in the array. In most cases, for small values of k, the dependence of our algorithms on the number n of elements is close to $\log n$, with only additive dependency on the number of imprecisions. In other words, these results smoothly interpolate beween parameter regimes where algorithms are as good as binary search and the unavoidable worst-case where linear search is best possible.

That said, why should one be interested in, e.g., almost tight bounds relative to the number of block moves that take A to a sorted array, as the bounds are far from binary search? The point is that only the total number of comparisons matter, and having a worse function that depends on a (in this case) much smaller parameter value can be favorable to having a much better function of a large parameter value. E.g., after a constant number of block swaps the parameters $k_{\max}$, k_{sum} etc. may have value $\Omega(n)$ and the guaranteed bound becomes trivial, while running the search algorithm for the case of few block swaps guarantees $\mathcal{O}(\sqrt{n})$ comparisons. Similarly, having tight bounds for the various parameters gives us the exact (worst-case) regime for the chosen parameter (in terms of n) where a sophisticated algorithm can outperform linear search, or even be as good as binary search.

Despite having already asymptotic tightness, it would be interesting to close the gaps between coefficients of dominant terms in upper and lower bounds for some of the cases. Another question would be to find a different restriction than $\text{pos}(e) = \text{rank}(e)$, i.e., the target being in the correct position relative to sorted order, that avoids degenerate lower bounds of $\Omega(n)$ queries for several parameters. A relaxation to allowing a target displacement of ℓ and giving cost in terms of n, k, and ℓ seems doable in most cases, but is unlikely to be particularly insightful. Finally, it seems interesting to study whether randomization could lead to improved algorithms for some of the cases. The analysis of randomized lower bounds requires entirely new adversarial strategies since the adversary must choose an instantiation without access to the random bits of the algorithm.

Acknowledgements. The authors are grateful to several reviewers for their helpful remarks regarding presentation and pertinent literature references.

References

[1] Arne Andersson, Torben Hagerup, Johan Håstad, and Ola Petersson. Tight bounds for searching a sorted array of strings. *SIAM J. Comput.*, 30(5):1552–1578, 2000.

[2] Stanislav Angelov, Keshav Kunal, and Andrew McGregor. Sorting and selection with random costs. In *Proceedings of the 8th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 48–59, 2008.

[3] Javed A. Aslam and Aditi Dhagat. Searching in the presence of linearly bounded errors. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 486–493, 1991.

[4] Jérémy Barbay and Gonzalo Navarro. On compressing permutations and adaptive sorting. *Theor. Comput. Sci.*, 513:109–123, 2013.

[5] Therese C. Biedl, Timothy M. Chan, Erik D. Demaine, Rudolf Fleischer, Mordecai J. Golin, James A. King, and J. Ian Munro. Fun-sort–or the chaos of unordered binary search. *Discrete Applied Mathematics*, 144(3):231–236, 2004.

[6] Biagio Bonasera, Emilio Ferrara, Giacomo Fiumara, Francesco Pagano, and Alessandro Provetti. Adaptive search over sorted sets. *J. Discrete Algorithms*, 30:128–133, 2015.

[7] Ryan S. Borgstrom and S. Rao Kosaraju. Comparison-based search in the presence of errors. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC)*, pages 130–136, 1993.

[8] Allan Borodin, Leonidas J. Guibas, Nancy A. Lynch, and Andrew Chi-Chih Yao. Efficient searching using partial ordering. *Inf. Process. Lett.*, 12(2):71–75, 1981.

[9] Gerth Stølting Brodal, Rolf Fagerberg, Irene Finocchi, Fabrizio Grandoni, Giuseppe F. Italiano, Allan Grønlund Jørgensen, Gabriel Moruz, and Thomas Mølhave. Optimal resilient dynamic dictionaries. In *Proceedings of the 15th Annual European Symposium on Algorithms (ESA)*, pages 347–358, 2007.

[10] F. Warren Burton and Gilbert N. Lewis. A robust variation of interpolation search. *Inf. Process. Lett.*, 10(4/5):198–201, 1980.

[11] Ferdinando Cicalese. *Fault-Tolerant Search Algorithms - Reliable Computation with Unreliable Information*. Springer, 2013.

[12] Aditi Dhagat, Peter Gacs, and Peter Winkler. On playing "twenty questions" with a liar. In *Proceedings of the 3rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 16–22, 1992.

[13] Vladimir Estivill-Castro and Derick Wood. A survey of adaptive sorting algorithms. *ACM Comput. Surv.*, 24(4):441–476, 1992.

[14] Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM Journal on Computing*, 23(5):1001–1018, 1994.

[15] Irene Finocchi, Fabrizio Grandoni, and Giuseppe F. Italiano. Optimal resilient sorting and searching in the presence of memory faults. *Theor. Comput. Sci.*, 410(44):4457–4470, 2009.

[16] Irene Finocchi and Giuseppe F. Italiano. Sorting and searching in faulty memories. *Algorithmica*, 52(3):309–332, 2008.

[17] Gianni Franceschini and Roberto Grossi. No sorting? better searching! *ACM Transactions on Algorithms*, 4(1), 2008.

[18] Michael L. Fredman. The number of tests required to search an unordered table. *Inf. Process. Lett.*, 87(2):85–88, 2003.

[19] Anupam Gupta and Amit Kumar. Sorting and selection with structured costs. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science, (FOCS)*, pages 416–425, 2001.

[20] Donald E. Knuth. *The Art of Computer Programming, Volume III: Sorting and Searching*. Addison-Wesley, 1973.

[21] Philip M. Long. Sorting and searching with a faulty comparison oracle. Technical Report UCSC-CRL-92-15, University of California at Santa Cruz, 1992.

[22] Harry G. Mairson. Average case lower bounds on the construction and searching of partial orders. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 303–311, 1985.

[23] Kurt Mehlhorn. Sorting presorted files. In *Proceedings of the 4th GI-Conference on Theoretical Computer Science*, pages 199–212, 1979.

[24] Kurt Mehlhorn. *Data Structures and Algorithms 1: Sorting and Searching*, volume 1 of *EATCS Monographs on Theoretical Computer Science*. Springer, 1984.

[25] S. Muthukrishnan. On optimal strategies for searching in the presence of errors. In *Proceedings of the 5th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 680–689, 1994.

[26] Andrzej Pelc. Searching with known error probability. *Theoretical Computer Science*, 63(2):185–202, 1989.

[27] Andrzej Pelc. Searching games with errors – fifty years of coping with liars. *Theoretical Computer Science*, 270(1-2):71–109, 2002.

[28] Ola Petersson and Alistair Moffat. A framework for adaptive sorting. *Discrete Applied Mathematics*, 59(2):153–179, 1995.

[29] Erez Petrank and Guy N. Rothblum. Selection from structured data sets. *Electronic Colloquium on Computational Complexity (ECCC)*, (85), 2004.

[30] Ronald L. Rivest, Albert R. Meyer, Daniel J. Kleitman, Karl Winklmann, and Joel Spencer. Coping with errors in binary search procedures. *Journal of Computer and System Sciences*, 20(3):396–404, 1980.

[31] Robert Sedgewick. *Algorithms in C++ - Parts 1-4: Fundamentals, Data Structures, Sorting, Searching*. Addison-Wesley-Longman, 1998.

[32] Andrew Chi-Chih Yao. Should tables be sorted? *J. ACM*, 28(3):615–628, 1981.

A. Relations between measures of array disorder

A.1. Equivalences

Proposition 5. $k_{\text{seq}} = k_{\text{mov}} = k_{\text{rep}}$.

Proof. $k_{\text{seq}} \leq k_{\text{mov}}$: If k_{mov} element moves lead to a sorted array then the at least $n - k_{\text{mov}}$ elements that are not moved must form a sorted subsequence.

$k_{\text{mov}} \leq k_{\text{seq}}$: The k_{seq} elements that are not part of any fixed ordered subsequence of length $n - k_{\text{seq}}$ can be moved to the correct positions in that sequence by k_{seq} element moves.

$k_{\text{seq}} \leq k_{\text{rep}}$: Assume that k_{rep} element replacements suffice to reach a sorted array. It follows that at least $n - k_{\text{rep}}$ elements are not replaced and their subsequence must be already sorted.

$k_{\text{rep}} \leq k_{\text{seq}}$: Let P a set of k_{seq} positions such that the subsequence S on the remaining $n - k_{\text{seq}}$ elements is sorted. One can replace the elements in P such that the entire array is sorted. □

Proposition 6. $k_{\text{aswap}} = k_{\text{inv}}$.

Proof. If there are any inversions then there is an adjacent inversion, which can be removed by a single swap of the adjacent elements; this lowers the number of inversions by one. On the other hand, each swap of adjacent elements affects only the relative ordering of these elements and, hence, removes at most one inversion. □

A.2. Relations by similarity of operations

The following relations hold because some number of operations in terms of the first measure can be used to implement one operation of the second measure.

Proposition 7. $k_{\text{swap}} \leq k_{\text{aswap}}$.

Proposition 8. $k_{\text{rbswap}} \leq k_{\text{swap}}$.

Proposition 9. $k_{\text{bswap}} \leq k_{\text{rbswap}}$.

Proposition 10. $k_{\text{bmov}} \leq k_{\text{mov}}$.

Proposition 11. $k_{\text{bswap}} \leq k_{\text{bmov}} \leq 2k_{\text{bswap}}$.

Proof. Any block move can be implemented by swapping the block with an empty block. Any block swap can be implemented by two block moves. □

Proposition 12. $k_{\text{rep}} \leq 2k_{\text{swap}}$.

Proof. A swap of two positions in the array can be implemented by two replacements. □

A.3. Further relations

Proposition 13. $k_{\text{inv}} \leq k_{\text{sum}} \leq 2k_{\text{inv}}$.

Proof. $k_{\text{inv}} \leq k_{\text{sum}}$: If at least one element is displaced then at least one element e occurs before $\text{rank}(e)$ and at least one element e' occurs after $\text{rank}(e')$. If any element e occurs before $\text{rank}(e)$ then there must be a subsequent element e', i.e., with $\text{pos}(e) < \text{pos}(e')$, that occurs after $\text{rank}(e')$. Consider a pair e and e' of elements with $\text{pos}(e) < \text{rank}(e)$ and $\text{pos}(e') > \text{rank}(e')$ such that all elements e'' between $\text{pos}(e)$ and $\text{pos}(e')$ have their correct position $\text{pos}(e'') = \text{rank}(e'')$.

(Possibly there are no such elements e'', but it should be clear that e_1 and e_2 can always be found if $k_{\text{sum}} > 0$.) It follows that $\text{rank}(e) \geq \text{pos}(e')$ since $\text{rank}(e) > \text{pos}(e)$ and the positions between $\text{pos}(e)$ and $\text{pos}(e')$ are already filled with non-displaced elements. Similarly, $\text{rank}(e') \leq \text{pos}(e)$. Consider the operation of swapping e and e': This would lower the total displacement by $2(\text{pos}(e') - \text{pos}(e))$ since $\text{rank}(e') \leq \text{pos}(e) < \text{pos}(e')$ and $\text{rank}(e) \geq \text{pos}(e') > \text{pos}(e)$. In terms of swaps of adjacent elements the swap of e and e' costs exactly $2(\text{pos}(e') - \text{pos}(e) - 1$. The relation follows since the number of adjacent swaps is at most the decrease in terms of total displacement.

$k_{\text{sum}} \leq 2k_{\text{inv}}$: Recall that $k_{\text{inv}} = k_{\text{aswap}}$. Swapping any two adjacent elements can lower the total displacement by at most two since the elements are moved a total of two positions. □

Proposition 14. $k_{\max} \leq k_{\text{inv}}$.

Proof. Recall that $k_{\text{inv}} = k_{\text{aswap}}$. Every swap of adjacent elements moves the elements by exactly one position each. Thus, the maximum displacement is lowered by at most one. □

Proposition 15. $k_{\text{rep}} \leq k_{\text{inv}}$.

Proof. Consider the first element, say x, in the sequence that is in an inversion (with some later element). It follows that elements preceding x are not larger than any later element. In particular, the directly preceding element, say y, must be strictly smaller than x and not exceed any later element. Now, if some later element is equal to y then replace x by y; else, replace it by an arbitrary value that is larger than y but smaller than any later element. Clearly, in both cases all inversions involving x are handled (at least one), proving the bound. (Note that the replacement rules ensure that arrays of unique numbers will retain this property. Setting x to the value of y is only done if needed, i.e., if that value already occurs at least one more time.) □

Proposition 16. $k_{\text{ainv}} \leq 2k_{\text{bswap}}$.

Proof. This can be proved by analyzing the three different types of block swaps: (i) block moves, i.e., swapping a nontrivial block with an empty block, (ii) swapping two nontrivial, nonadjacent blocks, (iii) swapping two nontrivial, adjacent blocks. Cases (i) and (iii) can be verified to only increase the number of adjacent inversions by at most two, else leading to a simple contradiction. For (ii), assume that we start with

$$\ldots, a, b, \ldots, c, d, \ldots, a', b', \ldots, c', d', \ldots$$

and swap $b, \ldots, c$ with $b', \ldots, c'$ to obtain

$$\ldots, a, b', \ldots, c', d, \ldots, a', b, \ldots, c, d', \ldots.$$

Note that only the eight pairs (a, b), (c, d), (a', b'), (c', d'), (a, b'), (c', d), (a', b), and (c, d') matter for upper-bounding the increase in number of adjacent inversions. If $a > b'$ and $a' > b$ then $a > b$ or $a' > b'$ must hold (depending on $b \geq b'$ or $b < b'$); in other words, if there are adjacent inversions at both (a, b') and (a', b) after the block swap then among (a, b) and (a', b') there was at least one adjacent inversion. Similarly, if $c' > d$ and $c > d'$ then $c > d$ or $c' > d'$ must hold, i.e., if we have inversions at both (c', d) and (c, d') then we had at least one adjacent inversion among (c, d) and (c', d'). Thus, the total of adjacent inversions increases by at most two with this block swap. □

Proposition 17. $k_{\text{ainv}} \leq k_{\text{seq}}$

Proof. Pick any k_{seq} positions P such that the subsequence S of the remaining $n-k_{\text{seq}}$ positions is sorted. Clearly, any adjacent inversions must be between elements of P or between an element of P and an element of S. Consider any subsequence $x, y_1, \ldots, y_p, z$ where $p \geq 1$ and $y_1, \ldots, y_p \in P$ and $x, z \in S$. If there is the maximum of $p+1$ adjacent inversions then it follows that $x > y_1 > \ldots > y_p > z$, violating that $x < z$ in the sorted subsequence S. Else, there are at most p adjacent inversions incident with $y_1, \ldots, y_p \in P$. Thus, overall, have at most k_{seq} adjacent inversions. □

A.4. Unboundedness results

Here we give pairs of measures such that the second can be unbounded, even if the first is constant. Each such relation is represented by a dashed red arc in Figure 1. Note that for any pair of parameters not connected with a directed path in the figure, unboundedness follows, because any bound would produce some path that contradicts an unboundedness relation.

The first proposition of this type covers the comparison of all parameters other than $k_{\max}$ with $k_{\max}$ since we showed that k_{ainv} is bounded whenever any other parameter (except $k_{\max}$) is bounded.

Proposition 18. *There exist arrays with* $k_{\max} = 1$ *and* $k_{\text{ainv}} = \Omega(n)$.

Proof. For given even integer n, consider the array $A = [2, 1, 4, 3, \ldots, n, n-1]$. Clearly, each element e is (exactly) one position away from $\text{rank}(e)$. However, we find that the array has $\frac{n}{2} = \Omega(n)$ adjacent inversions. □

Proposition 19. *There exist arrays with* $k_{\text{swap}} = 1$ *and* $k_{\max} = \Omega(n)$.

Proof. For given n, consider the array $A = [n, 2, 3, \ldots, n-2, n-1, 1]$. A single element swap suffices to reach a sorted array, but the maximum displacement is $n-1 = \Omega(n)$. □

Proposition 20. *There exist arrays with* $k_{\text{rep}} = 1$ *and* $k_{\max} = \Omega(n)$.

Proof. For given n, consider the array $A = [n+1, 2, 3, \ldots, n-2, n-1, n]$. A single element replacement, namely $n+1$ by 1, suffices to reach a sorted array, but the maximum displacement is $n-1 = \Omega(n)$. □

Proposition 21. *There exist arrays with* $k_{\text{rep}} = 1$ *and* $k_{\text{rbswap}} = \Omega(\log n)$.

Proof. We prove by induction on d that an array A of size $n = 4^d$ with a large element followed by an increasing sequence of smaller elements cannot be sorted with less than d restricted block swaps. The claim trivially holds for $d \in \{0, 1\}$. For $d > 1$, consider any shortest possible sequence of restricted block swaps to sort the array. Let A' be the array after executing the first swap only and let i be the index of $A[1]$ in A'. If $i \leq 3n/4$, we can apply induction on $A'[i, i+1, \ldots, i+n/4-1]$. If $i > 3n/4$, let j be the largest index with $A[j] \neq A'[j]$ and let j' be the index of $A[j]$ in A'. We can then apply induction on $A'[j', j'+1, \ldots, j'+n/4-1]$ to obtain the claimed bound. □

Proposition 22. *There exist arrays with* $k_{\text{rbswap}} = 1$ *and* $k_{\text{seq}} = \Omega(n)$.

Proof. For given even integer n, consider the array $A = [\frac{n}{2}, \frac{n}{2}+1, \ldots, n, 1, 2, \ldots, \frac{n}{2}-1]$. Clearly, A can be turned into a sorted array by a single restricted block swap, i.e., it has $k_{\text{rbswap}} = 1$. Its longest sorted subsequence, however, has length $\frac{n}{2}$, implying $k_{\text{seq}} = \frac{n}{2} = \Omega(n)$. □

Proposition 23. *There exist arrays with* $k_{\text{ainv}} = 1$ *and* $k_{\text{bswap}} = \Omega(n)$.

Proof. For a given even integer n, consider the array $A = [1, 3, 5, \ldots, n-1, 2, 4, \ldots, n]$. Clearly, A has a single adjacent inversion. To sort the array with block swaps, there needs to be a block swap that increases the distance between consecutive odd/even numbers. Each block swap can increase at most one such distance, hence we need at least $\frac{n}{2} - 1$ block swaps to sort the array. □